

Using Culture and Values to Support Flexible Coordination

This research was supported by Université de Montpellier 2 and Utrecht Universiteit.



SIKS Dissertation Series No. 2015-20

The research reported in this thesis has been carried out under the auspices of SIKS, the Dutch Research School for Information and Knowledge Systems.

© 2015 Loïs Vanhée

Printing: Ridderprint BV, the Netherlands

Based on the *L^AT_EX* template from Susan van den Braak

ISBN 978-94-6299-145-3

Using Culture and Values to Support Flexible Coordination

Het ondersteunen van flexibele coördinatie met behulp van cultuur en waarden

(met een samenvatting in het Nederlands)

Coordonner flexiblement en utilisant des cultures et des valeurs

(avec un résumé en français)

Proefschrift

ter verkrijging van de graad van doctor aan de Universiteit Utrecht op gezag van de rector magnificus, prof.dr. G.J. van der Zwaan, ingevolge het besluit van het college voor promoties in het openbaar te verdedigen op dinsdag 22 september 2015 des middags te 2.30 uur

door

Loïs Charles Bernard Cornil Vanhée

geboren op *17 september 1987 te Chatenay Malabry, Frankrijk*

Promotoren: Prof. dr. J.-J.C. Meyer
Prof. dr. J. Ferber

Copromotor: Dr. F.P.M. Dignum

This degree is awarded as part of a Joint Doctorate with the Université de Montpellier 2



Contents

1	Introduction	1
1.1	What's Your Problem?	2
1.2	From a Problem to a Goal	4
1.3	From Our Goal to Milestones: Research Questions	5
1.3.1	Towards Answering Research Question 1: <i>How Does Culture Influence Individual Decisions Regarding Coordination in Human Societies?</i>	6
1.3.2	Towards Answering Research Question 2: <i>How to Model the Influence of Culture on Coordination?</i>	8
1.3.3	Towards Answering the End-Goal Research Question: <i>How to Use Models of Human Culture as a Practical Tool for Supporting Flexible Coordination in Artificial Societies?</i>	8
1.4	Structuring This Thesis	9
1.5	Reading Notice	10
2	Background	13
2.1	Coordination	14
2.1.1	General Background: What is Coordination?	14
2.1.2	Coordination in Human Societies	16
2.1.3	Key Aspects to Consider When Coordinating Agents	19
2.1.4	Methods for Coordinating MASs	21
2.2	Culture	27
2.2.1	General Background: What Is Culture?	27
2.2.2	Characterizing the Visible Manifestations of the Influence of Culture	33
2.2.3	Core Mechanisms of Culture	36
2.2.4	Computational Models of Culture	38
2.3	Relating Culture & Coordination	42
2.3.1	General Background: Culture & Coordination	43

2.3.2	Manifestations of the Influence of Culture on Coordination	45
2.3.3	Key Mechanisms of Culture Involved in Situation of Coordination . .	46
2.3.4	Using Culture as a Means for Supporting Coordination	46
2.4	Chapter Summary	47
3	Towards a MAS-Compliant Conceptualization of Visible Influences of Culture on Coordination	51
3.1	Conceptualizing the Influence of Culture on Coordination Through Hofstede's Cultural Dimensions	53
3.1.1	Power Distance	53
3.1.2	Individualism	57
3.1.3	Masculinity	60
3.1.4	Uncertainty Avoidance	63
3.1.5	Long Term Orientation	66
3.2	Relating Value-Systems and Hofstede's Cultural Dimensions	69
3.3	Chapter Summary	71
4	Replicating the Influence of Culture on Coordination within MASs: from Theories to Models	75
4.1	Key Mechanisms of the Influence of Culture on Coordination in Human Societies	76
4.1.1	Determining the Key Mechanisms of Culture to Focus On	77
4.1.2	Further Introducing These Mechanisms	78
4.2	Modeling Value-Systems	81
4.2.1	Modeling Values	82
4.2.2	Modeling the Influence of Culture on Value-Systems	84
4.3	Modeling Culturally-Sensitive Agent Decision Processes	84
4.3.1	Memory Model	85
4.3.2	Decision Loop	86
4.4	Culturally-Sensitive Decision Aspects	89
4.4.1	Identity and Relationships	89
4.4.2	Norms	92
4.4.3	Uncertainty	94
4.4.4	Learning	96
4.5	Chapter Summary	97
5	Showing Evidence That Our Decision Architecture Can Replicate Properties of Human-Like Culture	101
5.1	Overview	102
5.1.1	Environmental and Coordination Background	103
5.1.2	Decision Aspects	103
5.2	Simulation Model	107
5.2.1	Environmental Model	107
5.2.2	Coordination Structure	109
5.2.3	Update Loop	110

5.3	Agent Decision Process	113
5.3.1	Memory Structure	114
5.3.2	Decision Loop	116
5.4	Comparing Our Model with Properties of Human Culture	120
5.4.1	Validating the Influence of Culture on Individual Decisions	121
5.4.2	Experimentation Setup	122
5.4.3	Validating the Influence of Culture on Collective Outcomes through Experiments	123
5.5	Additional Evidence	130
5.6	Chapter Summary	131
6	Explaining the Influence of Culture on Coordination from Individual Decisions to Collective Outcomes, Theory and Simulations	133
6.1	A Theory of the Influence of Culture on Coordination	135
6.2	Overview	138
6.2.1	Environment	138
6.2.2	Coordination Model	139
6.2.3	Agent Decision Processes	139
6.3	Baseline Simulation Model	140
6.3.1	Concepts & Objects	141
6.3.2	Coordination Model	143
6.3.3	Simulation Loop	144
6.3.4	Agent Decision Process	146
6.4	Simulation 1: Exploring the Baseline Model	149
6.4.1	Experiments	149
6.4.2	Back to Our Theory	151
6.5	Simulation 2: Mismatching Cultures— Cultures, Coherence of Interactions and Collective Outcomes	154
6.5.1	Experiments	154
6.5.2	Back to Our Theory	155
6.6	Simulation 3: Richer Culture—Relating Cultural Richness, Richer Interactions and Collective Outcomes	158
6.6.1	Simulation Model	158
6.6.2	Experiments	159
6.6.3	Back to Theory	161
6.7	Simulation 4: Restrictive Coordination Mechanisms— Relating Culture and Coordination Mechanisms	163
6.7.1	Simulation Model	163
6.7.2	Experiments	163
6.7.3	Back to Theory	164
6.8	Simulation 5: Changing Coordination— Relating Culture, Coordination Mechanisms, Recurrent Interactions and Collective Outcomes	164
6.8.1	Model	164
6.8.2	Experiments	165

6.8.3	Back to Theory	166
6.9	Further Evidence	168
6.10	Chapter Summary	169
7	Using Value-Systems for Supporting Coordination in Artificial Societies in Practical System Design	173
7.1	Answering Research Question 3.a: <i>When are Value-Systems Appropriate for Supporting Coordination?</i>	174
7.1.1	Theoretical Requirements for Applying Value-Systems	175
7.1.2	Relative Benefits: Specificities of the Support Provided by Value-Systems	176
7.1.3	Conclusion: Pragmatically Deciding When to Use Value-Systems for Supporting Coordination	180
7.2	Answering Research Question 3.b: <i>What Are the Practical Considerations for Designing a Culture to Support Flexible Coordination?</i>	180
7.2.1	Practical Aspects to Consider When Designing Values	181
7.2.2	Practical Aspects for Choosing the Values to Be Integrated	183
7.2.3	Practical Considerations for Combining Values	185
7.2.4	Practical Concerns for Integrating Value-Systems within Agent Decision Process	186
7.3	Conclusion: <i>How Can we Use Models of Human Culture as a Practical Tool to Support Flexible Coordination in Artificial Societies?</i>	191
8	Conclusion	193
8.1	Answering Research Questions	194
8.1.1	Answering Research Question 1: <i>How Does Culture Influence Coordination in Human Societies?</i>	194
8.1.2	Answering Research Question 2: <i>How to Model the Influence of Culture on Coordination?</i>	199
8.1.3	Answering Research Question 3: <i>How Can We Use Models of Human Culture as a Practical Tool to Support Coordination in Artificial Societies?</i>	201
8.2	Contribution	203
8.2.1	Contribution to Social Sciences	203
8.2.2	Contribution to the Agent Design Community	203
8.2.3	Contributions to the Simulation Community	204
8.2.4	Contribution to the MAS Design Community	204
8.3	Recommendations for Future Research	205
8.3.1	Social Science Theories	205
8.3.2	Modeling and Simulating Culture	207
8.3.3	Using Culture for Supporting Coordination	208
8.4	Closing Remarks	211
	Bibliography	213

Index	225
Acknowledgements, Dankwoord, Remerciements	225



1

Introduction

“Mais enfin, Loïs, un peu de bon sens !”
“Come on Loïs, have a bit of common sense!”

This is what one could frequently hear a few years ago, when my actions mismatched the expectations of my girlfriend. A memorable example was the time when I was caught washing the dishes with the red sponge. Before this time, I have always had only one sponge at home for washing the dishes. However, since we have started living together, there were two in the kitchen: a red one and a green one. For me, the closest was the best. For her, the green sponge was for washing the dishes while the red one was for cleaning the table. Thus, these two should not be mixed.

Her exclamation revealed me a coordination failure. We were working together in order to keep the house clean. However, our actions mismatched, leading to an undesirable outcome for both of us. I saw in this coordination failure a good opportunity for trying out my freshly-learned skills for coordinating agents from Multi-Agent Systems (MASs). I formulated it as a norm: “henceforth, the red sponge shall never be used for washing the dishes”. Hop, let’s rewash the dishes with the green sponge and that’s a problem solved!

Actually, this fix turned out to be quite limited. Less than ten minutes later, the same sentence could be heard when I was caught drying the dishes with the hand towel instead of the dish towel. Damn, another coordination failure and washing the dishes had to be done again. Right, things had not worked out as easily as expected. But let’s keep using MAS wisdom, this should eventually work out, right? Well, maybe not in this case: it did not take long before I was juggling with too many rules at the same time, while many, many more were to be found through trial-and-error.

I quickly discovered (on my own expense) the practical ineffectiveness of using traditional MASs techniques for coordinating actions in the “real world”. I do not say that these traditional techniques are useless: they can solve really well many coordination problems (e.g. setting a meeting, optimizing delivery of items, setting simple contracts). However, their use quickly becomes impractical when facing “real world” problems. These problems raise ineluctable annoying irregularities and unexpected situations (e.g. booking the defense room is a standard “set a meeting” coordination task with a small irregularity: available timeslots change over time). All of a sudden, traditional solutions are no longer appropriate (e.g. a doodle is insufficient). They either fail, particularly when handling

unexpected situations, or they require a lot of effort for handling all possible situations, being complex and costly to setup and to use (e.g. too many rules, learned through trial-and-error).

These methods lack *flexibility* for handling “real world” problems. Flexibility is achieved by being able to handle complex problems and a wide range of situations and goals. Several methods are appropriate for supporting *flexible coordination* (e.g. task-sharing, cooperation, teamwork). These methods mostly rely on agents for reasoning about coordination and generating coordination on the fly (e.g. the dirty red sponge has to be replaced, we reason and communicate about how to solve this issue).

However, these methods have certain limitations. One of the core problems lies in the different perspective of individuals when they make decisions (e.g. different interpretations about common goals, different criteria to optimize when planning, commitment). Coordination particularly offers an important room for subjectivity (e.g. what is important, following rules or instructions? Efficiency? Timeliness? Mutual support? Cleanliness?). These different perspectives can introduce misalignment of expectations, misunderstandings and, in turn, failures (e.g. “cleaning the dishes” leads to very different results when focusing on cleanliness or on efficiency).

This limitation is especially problematic when coordinating artificial agents, particularly in open systems where agents can be designed by different individuals who have themselves different perspectives. In this thesis, we aim to tackle this issue by proposing a method for achieving flexible coordination.

Nevertheless, in order to move towards this better world, we should first identify the exact problem.

1.1 What’s Your Problem?

The general problem we want to address in this thesis is the following:

Problem:

Available methods for achieving flexible coordination in MASs are limited. This limitation prevents using MASs for handling problems involving dynamic and complex environments or evolving goals.

The statement of this problem raises multiple conceptually-rich terms that deserve a brief introduction¹.

By coordination, we consider the task of influencing agents such that they achieve a goal in the environment². Agents are autonomous entities capable of making decisions based on their perception and internal memory. They can be influenced by coordinators in several ways (e.g. through messages, code). Agents are embedded within the environment. The environment is assumed to have a state that evolves by itself and that can be influenced by actions resulting from decisions of agents. Achieving a goal in the environment consists in leading the environment to a given state. In the context of MASs, coordi-

¹They are further detailed in Chapter 2.

²We do not assume that the entity performing this task is necessarily a human (e.g. the system designer), it can also be an agent within the system (e.g. a leader agent)

nators are generally system designers who coordinate software agents. This coordination generally only happens *a priori*, before the system is set to run.

By flexible coordination we mean being capable of coordinating agents such that they can handle a broad range of situations (states of the environment) and of goals. In particular, flexible coordination requires handling environments and goals that are complex and dynamic. The term “complex” is not formally defined, but it encompasses several categories of aspects, such as those with intrinsic conceptual complexity (e.g. building a spaceship), dense interactions between agents (e.g. logistics in a heavily loaded warehouse) and when actions have many direct and indirect consequences (e.g. puzzles). Dynamic environments can be related to quick evolution of the state of the environment with time, partly dependently on actions of agents (e.g. locating victims after a disaster recovery), evolving goals (e.g. stakeholders set on the fly goals for the system). As an illustration of complex and dynamic environments with evolving goals, we can consider systems for supporting humans (e.g. care robots in hospitals, disaster recovery robots) and tight logistics problems (e.g. on-demand delivery services).

Solving the problem of achieving flexible coordination is motivated by several incentives. First, this problem is important for the MAS community, which aims at using artificial agents for solving “real-world” complex and dynamic problems, particularly for helping humans who have evolving desires, as illustrated in the vision paper by Grosz (1996). Current limitations strongly restrict the range of applications for which MASs can be helpful. Currently, most applications are bound to coordination problems that are simple enough for limiting the spectrum of possible interactions (e.g. limited range of messages, limited conceptual basis) and, thus, limited uses of the system. Overcoming this limitation and offering more flexibility opens the door for using MASs in order to tackle a wider range of real-world problems. These problems can in particular involve and support humans. As a longer-term benefit, overcoming this limitation will enrich the autonomy of MASs: agents capable of flexibly coordinating are less dependent on human intervention in order to handle new situations and goals.

Second, this problem is particularly interesting for handling new coordination challenges raised by recent technologies. In the past, technology restricted the range of problems to relatively closed worlds (e.g. pure-software agents performing automatic trading, closed-world automated warehouses). Recent technological advances offer far richer means for integrating agents in open environments (e.g. affordable robots and computationally-powerful smartphones, houses, cars, fridges, ovens) and enabling rich interactions (e.g. cheap and efficient communication devices). Techniques such as ubiquitous computing offer richer-than-ever environments that can be exploited by intelligent agents.

Third, because clear progress can be expected in this direction. As a justification, MASs can be tightly related to human societies, which manage relatively well to achieve flexible coordination. Thus, flexible coordination is possible and, at least through replication of human societies, we can expect to obtain it within MASs.

Now that we have highlighted our problem and its importance, we should determine what should be done to tackle it.

1.2 From a Problem to a Goal

Before introducing the specific goal of this thesis, we first introduce the *abstract goal* we pursue. Then, we introduce a preliminary research for specifying the specific goal.

Abstract Goal:

We want to provide a practical method for better supporting flexible coordination in artificial societies.

This abstract goal raises two crucial aspects that influence the direction of the rest of this thesis. First, our method should be applicable to artificial societies³. To that extent, our method should be translatable into *computational models*. Second, our method should be practical. By practical, we mean that our method should be *practicable*, being usable in concrete cases (e.g. keeping acceptable design and computational costs). In addition, our method should provide *practical benefits* for coordinating agents, expanding coordination possibilities considering available coordination methods (e.g. by better handling problems that are difficult for available coordination techniques).

In order to achieve this abstract goal, we rely on social modeling. Briefly, this research method consists in replicating visible phenomena from human societies within MASs. More specifically for our case, we aim at replicating the aspects of human societies that make that humans manage to coordinate flexibly compared to artificial agents. This research method is relatively traditional when desiring to coordinate MASs (e.g. by integrating aspects of norms and organizations within MASs). This method relies on the strong parallel between human societies and MASs: many phenomena from human societies can be considered within MASs by relying on similar concepts. As a core benefit towards achieving our goal, this method allows relying on former social science research related to desired phenomenon. This former research indicates how the studied phenomenon can benefit to coordination. This former research also conceptualizes the phenomenon, providing a solid ground for building models. Finally, this former research provides empirical evidence, that can be used for validating models of the phenomenon. As a side bonus, our research method, through the replication of the human-like phenomenon, offers a possibility to better understand the replicated phenomenon (e.g. generative social sciences Epstein (1999)).

Relying on this research method requires determining which phenomenon to replicate. In other words, we need to be able of determining the underlying mechanisms that make that humans tend to coordinate with more flexibility than artificial agents. In order to determine this phenomenon, we performed a preliminary research for determining the most relevant social phenomenon.

In this direction, we considered a broad panel of social mechanisms. We investigated first phenomena arising from nature (e.g. ant societies, trust, power relationships). This range of phenomena shows strong limitations when aiming to reach flexibility: they are relatively static, failing to adapt quickly to new situations and goals. As an illustration, animal societies rarely and hardly engage in new activities (e.g. even if their survival is

³We do not exclude the integration of humans as agents within the MAS, but this integration is not our main focus.

threatened, ants will hardly become nomad and gorillas will hardly share food equally. No animal societies is likely to have the social means for coordinating in order to build space rockets).

Since mechanisms directly derived from nature are limited with regard to our goal, we consider coordination mechanisms arising from societies: in other words, we consider mechanisms arising from *culture*. As an illustration, imagine a group of randomly drawn individuals that are given a task requiring flexible coordination. If they are all picked from a different corner of the world, chances are high that they fail to coordinate, simply because they hardly understand each other's languages and behaviors. If they are picked from the same country, they are more likely to coordinate better even if minor disagreements and mismatches can be expected (e.g. creative vs. normative people). If people come from the same working team or family, one can strongly expect them to coordinate even better. The more their culture is shared, the better they coordinate.

In considering theories of culture in more details, culture is shown to be a key factor of the success of human societies for managing flexible coordination. Culture strongly influences the way individuals interpret the world and particularly social interactions (e.g. culture of strength and power, culture of harmony, culture of performance, culture of mutual care). Then, these various interpretations deeply influence the way individuals reason about what is coordination (e.g. when making decisions, shall I rely on leaders? Or comply with rules? Or care for the weak? Different answers obviously lead to different coordination outcomes) and create expectations about how others should reason and behave when coordinating (e.g. all expect others to comply with rules). When culture is shared, they share a social paradigm and thus manage to coordinate relatively well. When culture is not shared, "cultural clashes" can occur (e.g. misunderstandings, misalignment of expectations, miscommunication, distrust), easily leading to coordination failures.

This short analysis allows us to refine our goal: we aim to replicate the influence of culture in MASs in order to support flexible coordination. While not asserting that culture is the best or the only factor at play, culture still appears to be one of the core factors that enables or prevents flexible coordination in human societies. Given the importance of culture for achieving flexible coordination, we adopt the following goal for the rest of this thesis:

Our Specific Goal:

We want to understand better how models of human culture can be operationally used within artificial societies as a tool for supporting flexible coordination

1.3 From Our Goal to Milestones: Research Questions

Achieving our goal is not trivial. In order to determine how to act properly towards its resolution, we rely on a traditional technique: planning. We decompose our goal in a sequence of steps: accomplishing these steps should solve our goal.

Since our goal consists of understanding something, we express it through a research question. Then, we decompose this general question in multiple subquestions, which are more specific and simpler to answer, providing elements for answering the general

question. This process is then repeated for obtaining sufficiently specific questions that can be concisely answered.

The goal of this thesis can be expressed through the following research question:

End-Goal Research Question:⁴

How can we use models of human culture as a practical tool to support flexible coordination in artificial societies?

Answering this question raises additional questions to be investigated and understood first. In order to determine *how to use* models of human culture, we need to understand better *how to design* such a model of human culture. In turn, in order to design such a human-like model, we need to understand *how culture influences coordination in human societies⁵*.

In order to explore these questions in detail, we express them through the following two research questions:

Research Question 1:

How does culture influence coordination in human societies?

Research Question 2:

How to model the influence of culture on coordination?

These questions are decomposed in turn in subquestions. These decompositions are introduced in the following subsections.

1.3.1 Towards Answering Research Question 1: *How Does Culture Influence Individual Decisions Regarding Coordination in Human Societies?*

In order to answer this question, we consider the three canonical perspectives for considering MASs and their interrelations: individual decisions, collective outcomes and the environment.

In considering these three perspectives, the perspective of individual decisions is the only one that is necessary (but hardly sufficient) for modeling culture. Indeed, theories of culture explain that culture is embedded within individuals' minds, directly influencing their decisions. Visible influences of culture on collective and environmental perspectives are only side effects of the influence of culture on individual decisions. Thus, in order to answer the Research Question 1, we first investigate the following question:

Research Question 1.a:

How does culture influence individual decisions regarding coordination in human soci-

⁴This question is also referred to as Research Question 3. This question is later divided in two sub-questions, Research Question 3.a and Research Question 3.b

⁵Note that these questions match the standard steps for designing and using socio-inspired models: understanding the phenomenon in human societies; modeling this phenomenon; making use of this model.

eties?

This question is investigated in Chapter 2 and Chapter 3.

In spite of being theoretically unnecessary for modeling culture, collective and environmental perspectives still remain useful for building, validating and using models. These perspectives indicate interesting properties to model (e.g. in human societies, in culture C leader and subordinates are more tightly coupled than in culture C' . Does that also happen in our artificial societies?) and to validate (e.g. culture C supports higher performance in simpler environments than culture C'). In the case of using models, these perspectives indicate the type of collective outcomes and environmental properties to be expected from a culture. They can be used for purposefully introducing a certain desirable property within a MAS (e.g. using culture C for supporting coordination in a simple environment).

Taking our end-goal into consideration, we limit our exploration to the collective perspective. Indeed, this perspective relates relatively well with individual decisions and is particularly appropriate for considering coordination from a global perspective. The influence of culture on the environment is too indirect and appears not to be a core aspect that is at play for supporting coordination. Therefore, we prefer to leave it out from our research interest. Thus, this exploration can be expressed through the following question:

Research Question 1.b:

How does culture (indirectly) influence collective outcomes related to coordination in human societies?

This question is investigated in Chapter 2, Chapter 3 and Chapter 6.

In considering possible relationships between the investigated perspectives, the link between individual and collective is the most direct and interesting. This link is crucial for getting a clearer idea about how to replicate the influence of culture on collective outcomes, which is advocated by theories to result from the influence of culture on individual decisions. Studying this link raises the following research question:

Research Question 1.c:

How does the influence of culture on individual decisions lead to the influence of culture on collective outcomes in human societies?

This question is investigated in Chapter 6.

These three questions provide a better understanding of the core aspects to consider when building MASs models: how culture should influence agent decision processes, what behaviors are to be seen depending on culture, what collective behaviors should be seen and how these collective behaviors (should) result from the decisions and behaviors of individuals.

1.3.2 Towards Answering Research Question 2: *How to Model the Influence of Culture on Coordination?*

The answers provided by Research Question 1.a highlight how culture influences individual decisions in situation of coordination. As a spoiler, culture influences many aspects of individual decisions. Not all of these aspects are relevant for supporting coordination. For instance, the influence of culture on greeting (e.g. shaking hands vs. bowing) is at play in coordination situations, but does not influence coordination performance. Furthermore, certain aspects of culture are not appropriate for complex environments because they are too specific (e.g. culturally-sensitive specific rules for writing reports).

In order to build adequate models of culture for supporting coordination in complex environments, we need to determine the appropriate aspects of culture that can be used. These aspects are determined by the following question:

Research Question 2.a:

What are the key mechanisms of the influence of culture on coordination?

This question is investigated in Chapter 2 and Chapter 4.

Once the core aspects of the influence of culture on coordination have been determined, the question consists in finding the adequate technical solution for modeling them. This can be expressed by the question:

Research Question 2.b:

How can we build practical models of the key mechanisms of the influence of culture on coordination?

This question is investigated in Chapter 4. A validation of this model is proposed in Chapter 5 and Chapter 6.

These questions should provide enough initial information for building adequate models of culture and thus tackle the end-goal question.

1.3.3 Towards Answering the End-Goal Research Question: *How to Use Models of Human Culture as a Practical Tool for Supporting Flexible Coordination in Artificial Societies?*

Once we have a clearer idea about how to model the influence of culture on coordination, we can further investigate how to play with models of culture such that we operationally improve coordination in artificial societies.

As highlighted in our goal, two aspects are crucial for making good use of the models of culture in practice.

First, we need to determine in which coordination problems culture can provide a useful support for coordination. In particular, we should determine when culture *adds to* available methods for achieving coordination. This is explored through the following question:

Research Question 3.a:

When is culture appropriate for supporting coordination?

This question is investigated in Chapter 7.

Second, we need to determine the practicality of these models. We need to understand better the practical considerations of about using these models (e.g. design and computational matters). This aspect is investigated through the following question:

Research Question 3.b:

What are the practical considerations for designing a culture to support flexible coordination?

This question is investigated in Chapter 7. In addition, this chapter introduces several techniques for handling these practical concerns.

Now that we see better the steps required for achieving our goal, we can structure the thesis for completing them.

1.4 Structuring This Thesis

As a starter, Chapter 2 will introduce existing research about culture and coordination in MASs in general. This chapter particularly illustrates that available theories do not provide enough background for directly and sufficiently answering the research questions of this thesis.

Since Research Question 1 is a precondition for all other research questions, we will explore it first in Chapter 3. This chapter will be dedicated to expanding available theories from social sciences for relating culture and coordination. In particular, we will focus on introducing MAS-related concepts that are necessary for building our models. This chapter will answer Research Question 1.a and Research Question 1.b. Unfortunately, existing social science theories remain insufficient for answering Question 1.c. We will answer later this question through simulation. In order to do so, we need to answer first Research Question 2.

The answers to Research Question 1.a and Research Question 1.b will provide enough initial ground for tackling Research Question 2 in Chapter 4. This chapter will investigate the most relevant aspects of culture to be integrated for handling the goal of this thesis (Research Question 2.a). Then, based on this investigation, we will introduce an agent decision model that integrates these cultural aspects (Research Question 2.b). In order to ensure that we properly model cultural aspects, Chapter 5 will validate the conformity of this model against the findings from social science, using the findings from Research Question 1.b.

Chapter 6 will use the ground provided by this model for answering Research Question 1.c through simulation. In this chapter, we will use simulation for explaining and exemplifying how the known (indirect) influence of culture on collective outcomes results from the known (direct) influence of culture on individual decisions.

Finally, Chapter 7 will rely on the answers to Research Question 1 and Research Question 2 and the models and relationships explored throughout this thesis for answering the

End-Goal Research Question. This chapter will explore the range of coordination problems for which culture is appropriate to be used as a coordination technique in comparison with existing MAS coordination techniques (Research Question 3.a). This chapter will also investigate technical considerations involved when using culture for supporting coordination (Research Question 3.b) and proposes several solutions for handling them.

A final chapter will summarize the content of this thesis, highlight our contribution to the respective field of research and indicate future research paths.

Figure 1.1 introduces an overview of the links relating the various research questions, how each chapter relies on those answers and answers further questions. This overview provides a global idea of the whole argument of this thesis, where to find answers for each research question and how chapters support each other.

1.5 Reading Notice

This thesis faces two hard challenges that deeply impact (1) results to expect, and (2) how to conduct research. In turn, both aspects dramatically influence the content of our research. We want to raise this point now, in order to avoid readers having wrong expectations about what they will find in the rest of this thesis.

(1) We rely on the topics of culture, one of the softest, fuzziest and most complex concept from human-sciences, and of coordination, which is also a relatively soft, fuzzy and complex concept. Thus, we cannot rely on neat, formal and complete definitions, theorems and strong assertions indicating “things happen exactly this way” or “this implies that”. We only have specific studies indicating (context-sensitive) empirical evidence and correlations, and informal theories trying to gather pieces of evidence. Metaphorically, we are trying to picture an item while being in a dark room, where the only light comes out of a strainer.

Thus, readers should not be surprised to see little formalism or “exact answers” in comparison with standard Computer Science theses. We are not deriving theorems based on a formal definition. For now, we take a pragmatic perspective, we show that the (main known) aspects of culture can be replicated within MASs for supporting coordination. We aim at giving a general idea about how things can be done, without locking ourselves in a specific formalism. When culture and coordination will be better understood, this will be the time for proposing formal definitions and theorems.

(2) This thesis covers a very broad distance: we started from informal theories from social sciences, to models, to validation, to expanding available theories to, finally, using these theories for coordinating MASs. Each step had to be gone through in order to reach our research goal, while having only support from previous research for shortcutting this long path.

In order to cover this distance, this thesis relies on an *exploratory* approach. This approach is not intended to provide conclusive evidence (e.g. extensive validation, all-encompassing models, proofs). Instead, this approach aims at better understanding the problem, focusing on finding the initial key arguments and aspects for answering each research question, providing a comprehensive initial argument for the whole thesis.

Thus, readers should not expect extensive empirical validations or “proofs”, claiming the exactness of our findings. We rely on (controlled) extrapolation for building our theo-

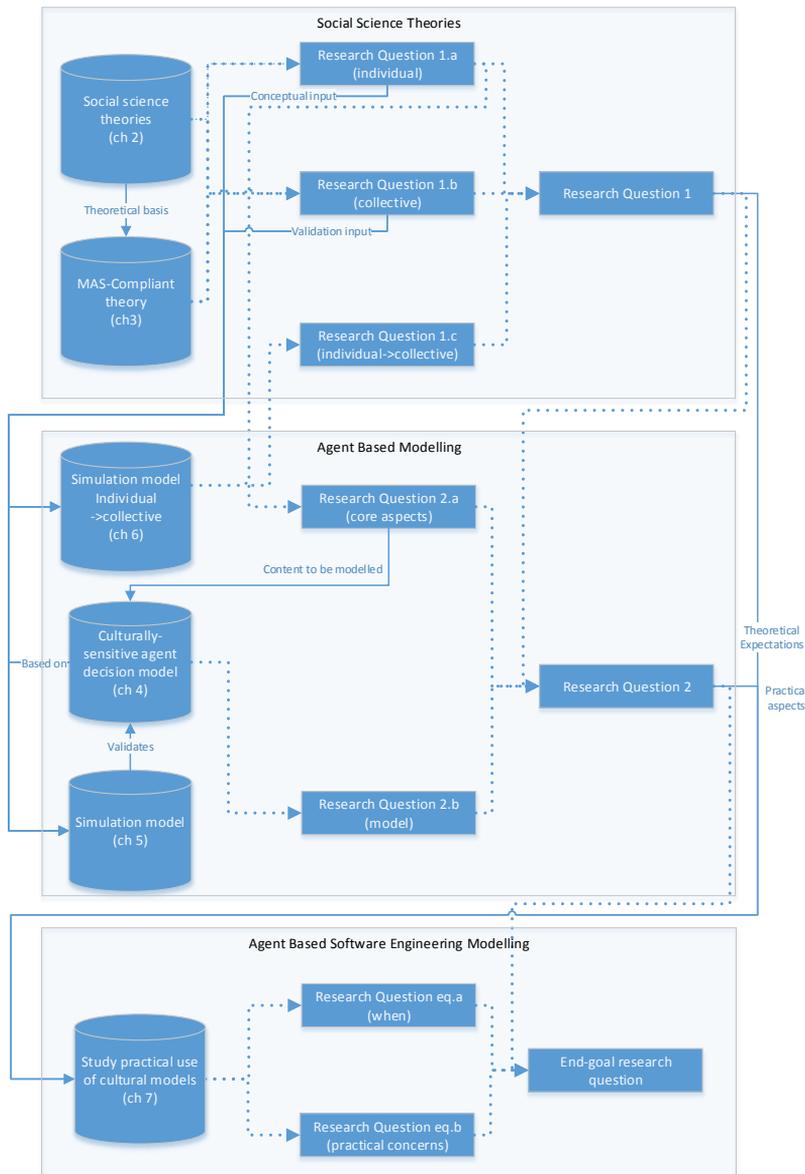


Figure 1.1 – Overview of the structure of this thesis. Dotted links mean that the start of the link helps explaining the end of the link.

ries and our models. Maybe future research will allow reconsidering or refining part of our arguments or models. Nevertheless, we were very careful to avoid wild speculation: we made sure to keep a strong link between our exploration and available theories and em-

pirical findings. Thus, while certain (minor, in our perspective) aspects may be arguable, the overall argument should stand.



2

Background

In this chapter, we explore existing work from former research that can help answering our research questions. Regarding our research questions, this exploration has two central purposes: (1) highlighting what is already known about these questions and (2) determining what remains to be explored in order to answer them.

The first purpose of this chapter aims at exploring former research towards answering our research questions. First, by directly answering our research questions or at least parts of them. Second, by providing a general background (e.g. definitions, theories, concepts, facts, perspectives) that can be used as a starting point for understanding topics raised by our research questions. In turn, this background provides a basis on which we can elaborate our own answers for our research questions. More generally, this exploration allows avoiding starting our research from scratch ¹.

The second purpose of this chapter consists in determining what is not yet covered and thus requires to be further explored in following chapters. Based on this, we can better determine which aspects we should focus on in order to answer our research questions and how we contribute to the advances of research in general.

In this thesis, we raise two topics that are analyzed in this chapter: coordination and culture. These two topics have been extensively explored in the past, thus a lot of previous research is available. While an exhaustive analysis of this literature would take too much space, we chose to introduce briefly each topic in two steps. First, we introduce a general background for understanding the core principles of each topic (e.g. examples, definition, key aspects). Then, we introduce former research that can be used for answering our research questions. After the independent introductions of coordination and culture, we explore their intersection in a third section.

This chapter is divided in three sections. First, Section 2.1 studies previous research about coordination. Then, Section 2.2 studies previous research about culture. Finally, Section 2.3 studies previous research describing the intersection of coordination and culture.

¹As Bernard de Chartres said: “nanos gigantum humeris insidentes”, we see further by sitting on shoulders of giants. In our context, these giants are the knowledge accumulated by previous research

2.1 Coordination

This section studies former research about coordination.

In a first section, we introduce a general background for better understanding the meaning of “coordination”. This term is particularly rich and is used in various contexts with slightly different meanings, thus it requires to be defined better in our context.

Then, in the following sections, we explore more specific aspects of coordination that are of interest in this thesis. First, we study existing research depicting concrete manifestations of coordination in human societies. This study aims at answering Research Question 1, which relates culture and coordination in human societies. Second, in order to achieve our goal as expressed by Research Question 3 we further investigate how to achieve coordination in artificial societies. In this direction, we first conceptualize what are the key aspects to consider when coordinating artificial societies. Second, since we aim at proposing a method for coordinating agents, namely, using culture for supporting flexible coordination, we study which methods are available for coordinating agents. In addition, we explore the key aspects that make that these methods promote flexible coordination (or not) and we introduce what is missing for supporting flexible coordination.

Section 2.1.1 introduces the general background for understanding coordination. Then, Section 2.1.2 introduces in more details concrete manifestations of coordination in human societies. Next, Section 2.1.3 introduces the core aspects involved in the task of coordinating. Finally, Section 2.1.4 introduces existing methods for coordinating agents in artificial societies.

2.1.1 General Background: What is Coordination?

Defining coordination is a challenging task: coordination is vague and has been intuitively used in many contexts without a formal definition. Furthermore, its meaning is often confused with related terms such as collaboration, cooperation, organization or collective action.

The current section aims at better capturing this term that is central for our thesis. In order to do so, we introduce a definition that can be used in our context for better describing what “coordination” and “to coordinate” mean. In addition, we briefly introduce terms with which coordination is often confused with (e.g. cooperation, collaboration) for avoiding ambiguities and reproducing similar confusion.

Defining Coordination

First of all, there is no exact definition of “coordination” that is consensually accepted by the research community. This term is both abstract and vague (e.g. one can hardly touch or measure coordination, in comparison with temperature or speed for instance). Plus, coordination has been informally and intuitively used in multiple domains, with similar yet different meanings (MASs, human societies, biology, physics).

To that extent, in the current section we propose a definition of coordination that can be practically applied in our context (MASs). This definition does not aim at being complete or exact (no sharp black & white definition). Instead, we aim at providing a comprehensive definition that is able to indicate relatively accurately whether coordination

occurs or not.

A General Definition: The most relevant general definition we found in literature was proposed by Malone and Crowston (1990). This definition is referred to as “a first general broad and common-sense” definition: “*Coordination is the act of working together harmoniously*”.

Malone and Crowston (1990) further specifies this definition. The concept of “working” is detailed in terms of actors that perform activities in order to achieve goals. These activities can raise *dependencies* between activities and actors. The term “harmoniously” informally qualifies how these interdependencies are solved (e.g. two actors trying to use the same resource at the same time). This definition provides important insight, but requires to be refined further for being more directly applied in the context of MASs.

Coordination in MASs: Actors are *agents* (e.g. humans, artificial agents) embedded in an *environment*. They *perceive* (or sense) this environment. They use this information and their internal memory for making *decisions* about which *action* they perform. Actions influence the environment. The environment can also evolve by itself.

Agents share (they sense and act in) a common environment. This sharing can create *dependencies* between agents (e.g. the actions of an agent can influence outcomes of other agents, agents can send messages to others). Agents are not necessarily aware of these dependencies.

In this context, observers can evaluate whether agents are coordinated by evaluating how “harmonious” are the actions of dependent agents. Harmony remains an informal term, but it can be related to objective measures (e.g. are agents trying to use the same resource at the same time? Are they bumping into each other?).

In order to avoid “false positives” (e.g. agents blindly running around but not bumping into each other by chance), we introduce an additional feature: *causality*. We assume that coordination does not happen only “by chance”. Instead, we assume that agents are necessarily coordinated *by something* (e.g. purposefully designed code, protocols, object in the environment).

From Coordination to Coordinating: In this thesis, we want to achieve more than recognizing coordination: we want to coordinate agents. In our context, coordinating is *goal-directed* (a goal is a property in the environment, such as “collecting a resource” or “extinguishing a fire”). Coordinators aim at influencing the resolution of dependencies towards both “harmonious resolutions” and achieving the goal.

In order to achieve this objective, coordinators can influence the MAS (e.g. designing agent code, imposing design constraints like norms, changing the environment), possibly under constraints (e.g. limited time for designing agents). In particular we assume the role of “human coordinators” (e.g. system or agent designers) that can act beforehand but not “at runtime”.

Given these definitions, this thesis investigates how to support coordination. More precisely, coordinators can influence agents by using a “culture” that they have to comply with.

Related Terms

Coordination is often used and confused with several related terms. In order to avoid introducing such confusion, we further introduce these terms and their difference from coordination. The terms we found out that are most frequently confused with “coordination”

are “organization”, “collaboration”, “cooperation” and “teamwork”.

Organization is the most difficult one as it encompasses multiple definitions. In its most general (and fuzzy) sense, organization denotes the emergence of structures within a system. We leave out this meaning in this thesis. Organization (Horling and Lesser (2005)) also refers to social structures (e.g. company). *Collaboration* is a form of intentional coordination where agents pursue a common goal, without considering plans for working together. *Cooperation* (Bratman (1992); Conte et al. (1991)) is a specific form of collaboration where agents establish collective plans for reaching the common goal. Finally, *teamwork* (Dunin-Keplicz and Verbrugge (2010)) is a specific form of cooperation where multiple agents achieve persistent goals. These terms are related to methods for coordinating MASs. These methods are further introduced in Section 2.1.4.

2.1.2 Coordination in Human Societies

This thesis aims at replicating the influence of culture on coordination as observed in human societies within artificial societies. To that extent, we need to understand better what coordination in human societies is. This understanding is used throughout this thesis, for answering Research Question 1 (what is the influence of culture on coordination in human societies) but also for modeling and validating models of coordination in human societies.

In this thesis, we particularly make use of a specific coordination technique: *organizations*. Organizations are by far the coordination technique that is the most studied in human societies (e.g. businesses, armies). Furthermore, available theories relating culture and coordination focus on organizations. This is the case with Hofstede et al. (2010a), the core theory of culture we rely on in this thesis, titled: “Cultures and Organizations: Software of the Mind”.

Former research about organizations depicts two main perspectives. First, the *formal perspective*. This perspective describes “top-down mechanical” organizations. This perspective abstracts away from humans, focusing on formal aspects of organizations (e.g. rules, responsibilities, communication, workflow). Second, the *informal perspective*. This perspective considers an orthogonal perspective: the “bottom-up human perspective”. In this perspective, organizations are just concepts given to humans who compose these organizations: other psychological aspects are at play (e.g. friendship, respect, charisma, boredom).

Each perspective is introduced in the following sections. Then a final section relates these perspectives by considering canonical organizational patterns.

Formal Perspective

The formal perspective focuses on formal mechanisms (e.g. rules) for coordinating individuals. In this thesis, we focus on the conceptualization provided by Morgenstern (1951), which is relatively complete and comprehensive from a MAS perspective.

From global perspective, organizations are social entities that pursue *organizational goals* (e.g. earning money, producing items, spreading knowledge, holding an event). These goals may be relatively vague (e.g. spreading knowledge in academia, helping people in hospitals). These goals drive the constitution of the organization and decisions made by its members.

Organizations can be represented by two main aspects: structures and responsibilities. Each *structure* describes a type of social interactions (e.g. delegation, monitoring) that is enabled by the organization. Structures relate *roles* with each other. Roles are to be enacted by entities (e.g. an individual, a machine). For instance “the production manager is connected to production workers through the delegation structure. Thus, the production manager can give orders to production workers”.

Multiple generic structures can be found. The *coordination structure* determines interaction in “standard” situations (e.g. the secretary warns the factory manager about an incoming production requests). The *delegation structure* indicates power or “leader-subordinate” relationships. The *control structure* indicates which roles monitor operations of other roles (e.g. verifying the quality of the products at the end of the production). The *information structure* indicates which information should be transferred from one role to the next (e.g. customer service manager reports the client satisfaction to the production manager). Finally, the *exception handling structure* indicates whom to contact if an unexpected situation occurs.

In addition to these structures, roles are given *responsibilities*. Responsibilities indicate the tasks and duties of individuals enacting the role (e.g. record requests from clients). Formally, roles can *operate* on the environment (e.g. producing an item) and receive and emit *signals* (e.g. send order, incoming calls from clients). Roles react to signals by operating and sending signals. Responsibilities react to these signals (e.g. if you receive a produce item request, operate accordingly then send a report signal). Occasionally, processing signals can raise *failures* (e.g. cannot process an order, failure in operation). In that case, failures are signal transferred through the exception handling structure.

This simple conceptualization can be enriched for handling more complex settings (e.g. managing resource, expertise). Organizations can change their structure over time.

Informal Organizations

The formal representation of organizations does not determine organizational behaviors. First because this representation is generally only partial (e.g. the responsibility “processing incoming orders” can be fulfilled in multiple ways) and rarely exactly complied with (e.g. individuals tend to create informal interactions which escape the formal structure). Thus, *informal* factors influence decisions made within organizations that indirectly impact coordination.

Multiple influences can be recognized. From a structural perspective, individuals create *informal relationships* (e.g. friendship) with each other (Waldstrom (2001)). These relationships escape the formal structure, which can lead to implicit structural changes (e.g. the secretary informally takes the responsibility of centralizing information).

From an individual perspective, individuals are rarely blindly committed to organizational goals. Instead, they generally aim at fulfilling their personal motives through their participation in the organization (e.g. getting a salary, meeting people, achieving something together). These motives are partly handled by formal organizations (e.g. salaries), but other informal aspects can impact decisions (e.g. deciding to stay late at work for helping a friend).

A last important informal aspect we particularly study in this thesis is status. Kemper (2011) explains that individuals can give each other *status*. An individual giving status to

another individual is more willing to act towards increasing the utility of the other individual (e.g. freely helping a nice colleague, supporting a charismatic leader). Kemper (2011) introduces the mechanisms by which individuals give status to each other and the importance they give to getting status. Statuses are argued to have a major informal influence on decisions made in the context of organizations.

Canonical Organizational Patterns

In this thesis, we aim at simulating organizations. In order to compare easily our simulations with reality, we simulate “canonical” well-studied organizational patterns. Mintzberg (1980) studied the grand categories of organizational structures that are used in human societies. Each organizational pattern enforces a unique paradigm for achieving coordination. This paradigm matches certain perspective for making individual decisions, occurring interactions, but also a general adequacy with specific types of environments.

Simple structures are coordinated through direct supervision. Their structure is one-level hierarchy. They have few leaders that manage many subordinates: they supervise, control, coordinate through delegation and handle failures. Leaders centralize information and make collective decisions. This organizational pattern performs well in dynamic environments (leaders can adapt their decisions on the fly) and in simple environments but not in complex ones (leaders are overloaded or out of expertise). Example: groceries.

Machine Bureaucracies (full bureaucracies from Hofstede et al. (2010a)) coordinate through standardization of work processes (i.e. extensive definition of responsibilities). Bureaucracies generally possess two structures. First, a vertical hierarchical structure relates leaders and subordinates. This structure merges delegation, monitoring and failure-handling. Second, a horizontal coordination structure relates roles that perform successive operations for standard signals. In this type of organizations, leaders coordinate by creating responsibilities for subordinates (standardizing through rules). Because of this standardization, this organizational structure handles well simple and repetitive tasks (e.g. mass production factories). This standardization prevents handling dynamic environments and centralization of standardization around leaders prevents handling too complex environments, otherwise leaders are overloaded.

Professional Bureaucracies rely on standardization of expertise as a coordination mechanism. Individuals are given “responsibilities” through heavy training to professional standard. Individuals tend to focus on completing tasks within their field of expertise in relative autonomy. This organizational pattern is particularly adequate for handling stable environments (expertise can hardly be changed on the fly) but experts can handle complex problems. Examples: hospitals and universities.

Adhocracies coordinate in relying on its informal structure (e.g. mutual adjustments, local adaptations). This form of organization minimizes the formal structure and supports informal contacts. In practice, individuals tend to generate solutions on the fly in teams of matching expertise for processing incoming signals. This organizational pattern can handle relatively well complex and dynamic environments by gathering the adequate groups of experts (e.g. software development).

Finally, *divisionalized forms* are macro-organizations composed of multiple sub-organizations. Sub-organizations tend to share similar standards but act in relative independence. This form is considered as an initial phase before bureaucracy.

2.1.3 Key Aspects to Consider When Coordinating Agents

Coordinating agent societies is complex, many aspects are to be considered for achieving this goal. This section introduces the core aspects of coordination that have been recognized by previous research.

In order to achieve this goal, we explore three aspects. First, we introduce a general background for capturing the core challenges of coordination in practical settings. Second, we introduce the core technical aspects to consider for determining which kind of solution to apply for solving a given problem and how to evaluate these solutions.

Challenges Raised by Coordination

This thesis considers a practical software engineering perspective and inherent constraints. From an engineering standpoint, coordinating is costly. Coordinating involves influencing agents (e.g. designing specifications and software, modifying the environment) and this influence is costly (e.g. time, money). Further on this track, even gathering information about the coordination problem or making decisions about how to coordinate is costly.

These costs entail pragmatic decisions from coordinators: they have to decide efficiently how to coordinate efficiently agents! In practice, this means discarding optimality for efficiency and sufficiency (achieving coordination objective). This pragmatism supports the use of well-oiled methods, which are certainly suboptimal, but are sufficient for achieving goals (e.g. designing this mechanism will cost relatively large amount of time and slower my agents, but this mechanism will make sure that objectives are achieved).

These constraints raise an important property with regard to the aim of this thesis. *The solution proposed in this thesis should add to existing solutions.* Our method should not only be useful for coordinating agents, but it should be more useful than available methods at least for certain problems. Otherwise, our method will be dropped for pragmatically other more cost-efficient ones.

Problem and Solution Features

Coordinating requires to consider deeply the coordination problem to tackle. In order to speed up this investigation and thus make efficient decisions towards coordinating agents, certain key aspects can be considered. These key aspects indicate the features of the problem to handle. These features can then be used for efficiently determining what to focus on when coordinating (e.g. watching out for unexpected situations) or for selecting an appropriate available technique and method for coordinating agents (e.g. need smart agents, a few norms can handle most of the problems).

In this direction Dignum and Dignum (2010) introduce four core aspects that are used in practice for classifying coordination problems: environmental features, social characteristics, autonomy characteristics and coordination characteristics. To these aspects, we integrate a fifth one, which is important with regard to the goal of this thesis: performance.

Environmental Features: The environment obviously influence the coordination problem and according solutions. More specifically, two main aspects are used for classifying environments: uncertainty and complexity.

Uncertainty evaluates the difficulty to predict environmental behavior. Uncertain environments have multiple causes such as quickly evolving environments, underlying randomness and complexity for making predictions. In considering coordination, uncertainty restricts possible assumptions that can be made about the environment.

Complexity evaluates the difficulty of seizing the consequences of decisions. Complex environments can be caused by rich dynamics (e.g. long-term consequences of decisions, decisions that impact many aspects) and tight interactions between agents. In considering coordination, environmental complexity requires richer agent decision capabilities or relying on heuristics.

Social Characteristics: solutions can be required to match certain characteristics. Dignum and Dignum (2010) highlight two core types of characteristics: system openness and coordination explicitness.

System openness: In closed systems, the set of agents is fixed and defined from scratch. New agents cannot be introduced. In open systems, external agents can relatively freely join and leave the environment (e.g. an online auction system). From a coordination perspective, closed systems can be tuned better for specific agents, but they cannot be expanded with new agents.

Explicit/Implicit Coordination: Implicit coordination solutions embed coordination within agent decision processes. Explicit coordination solutions are “public”. They can be related to specifications about coordination provided by the system designer or implemented coordination mechanisms, externally to agent decision process (e.g. norms to be complied with). From a coordination perspective, implicit coordination is adequate for allowing tightly embedded interactions (e.g. ant systems). Explicit coordination requires a relative separation between agent decision processes and coordination mechanisms.

Autonomy Characteristics: agents and societies can be more or less autonomous, bounding or not the range of achievable individual and social outcomes.

Goal Autonomy: This aspect indicates whether agents are seen as autonomous entities joining social behavior (pursuing their own goals) or as entities dependent on the society for behaving (pursuing social goals). From a coordination perspective, systems with goal autonomous agents tend to be more complex and difficult to predict, but they enhance the adaptability of the system.

Social Configuration: This aspect indicates whether the structure of society is fixed or can evolve with time (e.g. creation of new norms or roles). From a coordination perspective, static configuration makes the system more predictable but it limits its adaptability.

Coordination Characteristics: This aspect indicates whether the concepts for understanding social processes are static or can evolve with time and how much system control can be performed. In some sense, this aspect determines whether system owners can always understand and control their system.

Emergence: Emergence indicates how far the concepts for explaining the dynamics of a system can evolve, leading in particular to new social properties. From a coordination perspective, emergence can introduce desirable properties (e.g. ants finding optimal path), but allows the occurrence of unexpected behaviors, which may not be desirable (e.g. for high-security banking).

Locus of Control: The locus of control indicates how much information about the system can be obtained. Certain problems can be handled by keeping local information while others require a global overview of the system (e.g. need to be able to find agents with crit-

ical information). From a coordination perspective, the need for more control can limit the range of possible applications.

Performance: This aspect indicates the kind of results that a given solution can achieve. Three main indicators of performance are used for depicting the performance of a MASs: efficiency, flexibility and robustness. Unlike previous aspects, these three indicators tend to conflict with each other in practical aspects (more efficiency generally lowers flexibility and robustness).

Efficiency evaluates the quality of the resolution of standard cases. This evaluation is the most used and can be related to numerous concrete measures (e.g. minimal or average time for completing tasks, costs).

Flexibility evaluates to the range of problems that can be handled by a given solution for achieving coordination. Flexibility is difficult to measure formally, but can be related to the capability of the system to handle a broad range of goals (e.g. picking moving items, picking heavy items, picking multiple items) and environments (e.g. pick items in the sand, in the grass, in the factory). Flexibility also can be observed when the system handles correctly unexpected coordination problems. *In this thesis, we aim at proposing solutions for achieving high flexibility.*

Robustness evaluates how many “adversarial” events are required for preventing the system to achieve its goal. In general, considered adversarial events are failures (e.g. crashing agents) and overloads (e.g. many tasks requiring the same expertise arise at once).

2.1.4 Methods for Coordinating MASs

As indicated in Section 2.1.3, coordinating agent societies is a challenging task that has to be done within cost constraints. In order to limit costs, several methods have been proposed for helping coordinators to make quick, yet accurate decisions for supporting coordination. These methods consist of formalisms and techniques that can be concretely implemented for handling coordination problems. *In this thesis, we aim at proposing such a method: using culture for supporting coordination.* This aspect is particularly studied through Research Question 3.

Answering this question requires studying existing coordination methods. Indeed, as indicated in Section 2.1.3, our method should add to existing coordination methods. To that extent, we need to capture better the range available coordination techniques, what kind of problems they can handle and what are their pitfalls.

A first study of existing research indicates that many methods are available. In order to get a general overview, we depict the main categories of coordination methods. Given the goal of this thesis (“flexible coordination”), these categories of methods are based on and sorted by the flexibility they provide. A first subsection introduces these categories of methods, ordered by increasing flexibility.

Based on this list of methods, we can better understand what makes that a method is flexible or not. This understanding indicates the kind of factors we should look for in our own method. Furthermore, these factors can be used for investigating what is missing in the available literature for achieving flexible coordination.

Coordination Methods

This first category introduces the least flexible methods for coordinating agents. In spite of being relatively inflexible, these techniques are not worthless: they are appropriate for handling certain types of problems (e.g. very high efficiency, need simple agents). Thus, these methods may be used in combination with ours. Furthermore, their study can help understanding where (in)flexibility comes from.

Totalitarian DEC-POMDP Approaches: DEC-POMDPs Oliehoek (2012) (or DEC-MDPs Bernstein et al. (2002) if the environment is fully observable) formalize the dynamics of environments that are: stochastic (environmental evolution is not deterministic, random factors are at play), influenced by multiple agents (agents make actions influencing environmental evolution) and partially visible by agents (agents partly perceive the state of the environment). This formalism is used by optimization algorithms that compute a decision function for agents leading to optimal outcomes, given a function that relates the state of the environment to collective rewards. Thus, these algorithms fully determine the behavior of agents.

Considering its main features, DEC-POMDPs support extremely high efficiency but low flexibility and robustness (only the standard environment is considered). Furthermore, they fully determine agent decision process, making it hard to be connected to other coordination mechanisms. This method is thus not suitable with regard to this thesis.

Nature-Inspired Approaches: These methods focus on copying coordination mechanisms from artificial societies. These mechanisms are generally adequate for handling simple dependences and require limited cognitive resource (e.g. foraging ants from Panait and Luke (2004), flocking behavior from Reynolds (1987)).

Considering the main features, this form of coordination is generally relatively robust (agents can “die” without destroying the society) and simple. Nevertheless, these mechanisms are adequate for handling limited interactions that arise in natural environments. They can hardly be altered or used out of context for handling wider goals. Thus, they support low flexibility and can hardly be combined with other techniques. Thus, we leave out this coordination technique from this thesis.

Stygmergy: This method relies on the environment for coordinating agents (e.g. pheromones, semaphores, blackboards, sirens, round-robins). Omicini et al. (2004) propose a relatively generic formalism for introducing artefacts within the environments for supporting coordination.

Stygmergy is relatively useful for supporting or enforcing coordination in relatively simple and specific coordination problems (e.g. exchanging messages, raising signals, queuing). Nevertheless, stygmergy is relatively inflexible: artefacts cannot be dynamically changed without introducing costs. Furthermore, they are relatively passive, supporting only limited uses. Nevertheless, stygmergy provides an interesting support for simple and recurrent coordination problems (e.g. a blackboard (only) solves the problem of sharing a common memory, but this solution can be applied in many settings).

Coordinating with Specific Rules: Protocols, Norms and Organizations: These methods introduce frameworks for coordinating agents with rules (i.e. rule-based systems). These three methods combine relatively well with each other and share relatively similar features.

Protocols (Marzougui and Barkaoui (2013)) aim at ruling interactions. In general, pro-

protocols are represented by a state machine, where states indicate the situation of the interaction and transitions consist of actions of agents (e.g. communications). Agents are (partially) aware of the state of the protocol. The state of the protocol indicates which action can be triggered and its consequences. For instance, the Contract-Net Protocol formalizes interactions between a contractor who proposes a task and agents that can accept and process it.

Norms are inspired by human law systems (Bicchieri (2006); Castelfranchi et al. (2000); Dignum (1999)). Norms indicate general restrictions on the behavior of agents. When coordinating, these rules can then be used for preventing behavior that can damage collective outcome. The main framework for representing norms is deontic logic (Von Wright (1951)), which describes norms using obligations, permissions and prohibitions.

Organizations are inspired by formal human organizations as introduced in Section 2.1.2. At the core, organizations aim at formalizing responsibilities of agents and efficiently driving the flow of communication between them. Many formalisms have been developed for representing organizations such as *Moise*⁺ by Sichman et al. (2002) or OperA Aldewereld and Dignum (2010).

These three mechanisms can be used for handling coordination as human societies do. In comparison with previous methods, rule-based systems can be adapted for handling a wide range of problems. Nevertheless, they remain relatively specific. They can achieve relatively high efficiency, but overruling tends to prevent flexibility and robustness. Furthermore, they require handling all possible cases, thus can become too expensive to setup when the environment becomes too complex (e.g. overregulation as depicted by Castelfranchi (2000)).

Abstract Rule-Based Methods: In comparison with the previous methods, these methods aim at coordinating agents with relatively abstract rules. Instead of rules of the form “if you see a fire, send a message “warning:fire” to the firefighting team leader”, rules can have the form “dangers should be reported to competent authorities”. These latter rules are subject to interpretation and can be implemented in many ways (e.g. no definition of “danger”). Nevertheless, they allow expressing more concisely the general direction given to coordination. Thus, rules can be used for handling more complex environments and more situations at once (e.g. the rule is applicable for handling flooding and wild animals). Nevertheless, they offer less certainty about behaviors of agents and prevent formalizing tight interactions.

Internalizing Formal Coordination Method: This method consists in expanding traditional coordination techniques with mechanisms for updating them within the MAS, at runtime (e.g. expansion of the norm-base with rules for ensuring fairness when a resource becomes rare). In general, these mechanisms are implemented by adaptable norm bases, organizations and protocols (Boella et al. (2006); Hubner et al. (2004)). Often, an agent is in charge of managing the structure, but this can also be achieved with an artefact or an automated mechanism, with more limited intelligence.

Learning: (Panait and Luke (2005); Stone and Veloso (2000)) This method consists in rewarding agents when they perform desirable behaviors (*a posteriori*). Agents can use these rewards for determining whether their behaviors were adequate and adapting them accordingly. This method is useful for adapting individual decisions to each other and to evolving environments. Nevertheless, this method has numerous drawbacks: it is relatively slow, non-proactive (adaptation to failures that could have been avoided by ade-

quate preparation). Furthermore, learning can lead to undesirable behaviors (e.g. cyclic mutual adaptations, competition). In this thesis, we leave out this method.

Coordinating By Showing Rewards: This method consists in coordinating by influencing decisions of agents by allocating them rewards depending on collective outcomes (e.g. high reward if the product is finished on time). This coordination method is particularly studied through mechanism design or implementation theory, as explored by Maskin (2002). This framework relates well with Game Theory (Binmore (2008)) which is used for determining agent decisions and collective outcomes. Agents use the information about the rewards that other agents will get in order to create expectations about decisions of other agents. In turn, agents use these expectations for determining the most rewarding decision they should pick.

Briefly, while having numerous differences with our methods, this method is the closest from the coordination method proposed in this thesis. Our coordination technique relies on shared values supported by culture. These values indicate what is “good” and can be related to rewards. Furthermore these “rewards” are shared and thus agents can create expectations about decisions of others, thus making matching decisions. Nevertheless, our work strongly differs from available frameworks: our representations of values and their influence on decision mechanisms have little in common with game-theoretical frameworks presented in this section.

Integrating Agents Capable of Reasoning about Coordination: The last and most flexible solution consists in integrating agents capable of reasoning about coordination. In a certain sense, agents are given the responsibility of and the means for achieving coordination when facing a concrete problem.

Numerous techniques are available for doing so. The most basic techniques consist in introducing a coordinator within the system (e.g. a leader agent) in charge of coordinating agents that cannot reason about coordination (e.g. allocating independent tasks to subordinates).

This technique can be enriched by including multiple agents capable of setting up coordination. For instance, task or result sharing from Durfee (2006); Yeoh and Yokoo (2012) propose to let agents decompose tasks into subtasks that are then proposed to other agents. Then, agents can either further decompose this task or perform it. This form of decomposition leads to a *hierarchical decomposition*. This general framework raises numerous technical questions, which are investigated by related research (e.g. checking the consistency of the plan, transferring information when executing the plan, managing failures and changes in the plan, optimizing the parallel execution of the plan).

Finally, in an even more flexible setting, agents are all capable of reasoning about coordination. Thus, they *mutually* try to coordinate (with) each other, without any prior hierarchy for determining who coordinates who. These mutual influences can raise high complexity that frameworks aim at tackling. In this direction *cooperation* (Bratman (1992); Levesque et al. (1990); Wooldridge and Jennings (1999)) supports interaction frameworks where agents aim at agreeing on shared goals and collective plans for achieving these shared goals. Likewise, *teamwork* (Dunin-Keplicz and Verbrugge (2010); Tuomela (2000)) expands this formalism for considering multiple agents working together and recurrent interactions.

Factors of Flexible Coordination

In comparatively studying all these methods, we can investigate the key factors that make that a coordination method leads to high flexibility or not. These factors provide a conceptual basis for determining whether our method supports flexible coordination or not.

At the core, the main difference lies in whether the coordinator tries to *achieve* coordination or to *support* coordination.

In the first case, the coordinator tries to “force” harmonious resolutions of dependences (e.g. through rules, through the environment). Agents are just required to comply with the influences of the coordinator in order to achieve coordination. This form of coordination can thus allow relying on relatively simple agents who aim at complying with instructions and the coordinator can tightly interleave decisions of actions, leading to high efficiency or robustness. However, the coordinator has to coordinate *a priori*, before facing the concrete problem. Thus, coordination can fail if the environment was not the one expected by the coordinator or if new goals are set. Furthermore, system designers have to predict and handle all possible situations that all agents can face and react accordingly. Therefore, this form of coordination becomes quickly impractical when the environment is too complex (e.g. problems of overruling depicted by Castelfranchi (2000)). All these aspects make that “achieving” coordination leads to poor flexibility.

In the second case, the coordinator helps agents to coordinate in indicating general guidelines about how to interact. These guidelines should be “instantiable” in a wide range of problems. Then, agents use these guidelines for making intelligent action. They use these guidelines for making decisions towards achieving tight coordination. Thus, tight coordination is performed *within the system* and *on the fly*, when facing a concrete coordination problem. In considering performance, instantiating these general guidelines induce costs, preventing optimal efficiency. Nevertheless, this potential for adaptation allows to handle a wider range of environments (higher flexibility) while keeping relatively limited design costs for handling complex problems.

How to Support Better Flexible Coordination?

Available methods for supporting flexible coordination can be improved. In spite of “good intentions” agents can still fail to coordinate. Somehow these frameworks are still insufficient for enabling agents to create adequate expectations about each other, leading to failures. This problem is particularly disturbing in open systems, when coordinating with incoming agents that have different designers.

As an explanation, at the core, these methods rely on abstract and generic concepts for reasoning about coordination (e.g. cooperation, commitment, delegation, planning). Implementing these abstract concepts is subject to interpretation. Thus, multiple implementations can be proposed and are arguably acceptable (e.g. high commitment is important for not changing goals too often and achieving nothing, low commitment is important for avoiding wasting time on unfeasible goals and achieving nothing). However, these multiple interpretations and implementations can lead to mismatches (e.g. an agent is very committed to common plans, but not other). In turn, these mismatches can lead to misexpectations and thus to failures (e.g. both agents agree on working together a plan before splitting, the situation evolves and the plan is harder to achieve, an agent keeps trying to achieve a goal while the other dropped it, leading to failure).

More fundamentally, this observation highlights a problem of current methods for achieving coordination. Agents are in charge of internally coordinating the system. Thus, they are in charge of “bringing harmony” within the system. However, nothing tells them what is this “harmony” in general. Furthermore, they may have a very different idea of “harmony” with regard to other agents. With this in mind, failures which occur when trying to coordinate with flexible coordination techniques are relatively unsurprising.

As a hint, this thesis handles this issue with culture. Somehow, culture, through values, provide general guidance for specifying a collectively shared idea about what is “harmony” (e.g. when to be committed, when to drop plans, which aspects to prioritize when planning). This specification can then be used by cooperative agents for purposefully “working together” towards the same direction but also for creating strong mutual expectations about decisions of other agents.

Summary

In this section, we first introduced a general background for capturing what coordination means in the context of MASs. Briefly, coordination consists in influencing the system such that agent dependences are “harmoniously” solved. Then, we introduced more specific aspects of coordination that are used in this thesis.

Towards conceptualizing coordination and building simulation models, we first presented what coordination in human societies is, particularly through organizations. Organizations can be seen from two perspectives. The formal perspective introduces organizations as a structure for coordinating individual decisions, interactions communication and so on. The informal perspective introduces organizations as places where human-to-human informal business occurs, implying numerous irrational and psychological influences.

Then, we explored the core aspects to take in consideration when aiming at coordinating agents. This exploration indicates the challenges of coordinating in practical settings, pushing coordinators towards making *pragmatic* decisions (e.g. discarding optimality). In addition, we introduced core aspects of problems and solutions to consider when determining which coordination method to focus on: environmental features (uncertainty, complexity), social characteristics (system openness, explicit vs. implicit coordination), autonomy characteristics (goal autonomy, social configuration), coordination characteristics (e.g. emergence, locus of control) and performance (efficiency, flexibility, robustness).

Finally, we introduced available mainstream methods for achieving or supporting coordination. We explored which methods can be combined with additional methods. Through this study, we identified better the core aspects for achieving flexible coordination: coordinators should provide abstract means for supporting coordination that can be adapted or implemented by the system (e.g. agents) on the fly, when facing concrete problems. Finally, we introduced a core fundamental limit of methods for achieving flexible coordination: while they have abstract means for setting up coordination, current methods do not provide means for helping agents to agree on concrete forms of coordination to pursue, leading to misalignment of expectations and to failures.

2.2 Culture

This section studies existing research about the topic of culture.

As a first step, we first introduce a comprehensive general background for understanding better the fundamentals of culture. Furthermore, based on this general background, we determine which aspects of culture are relevant with regard to this thesis and which aspects are left out of our focus.

Then, we introduce in more details former research that is to be used in this thesis. First, we study the visible *manifestations* of the influence of culture. This study provides an initial answer to Research Question 1 (how culture influences coordination).

Then, we introduce the known *mechanisms* of culture. These mechanisms capture the central aspects that are (as directly as possible) influenced by culture that, in turn, lead to concrete cultural manifestations. The study of these mechanisms aims at providing the background for answering Research Question 2.a (what are the core mechanisms of culture).

Finally, we study available computational models of culture. This study aims at determining whether there are available models that we can be inspired by (or not) in this thesis. The study of these models aims at providing an initial answer to Research Question 2.b (how can we model the key mechanisms of the influence of culture).

Research Question 3 is out of the scope of this section. This question is studied in Section 2.3, which studies the intersection of culture and coordination.

Section 2.2.1 introduces a general background for conceptualizing culture. Section 2.2.2 introduces the core manifestations of culture. Section 2.2.3 introduces the core mechanisms of culture. Finally, Section 2.2.4 introduces the types of available computational models for representing culture.

2.2.1 General Background: What Is Culture?

Culture is manifested in many places: in what people say, what they think, how they behave, in music, in architecture, in societies. Specific manifestations of culture can easily be acknowledged when facing differences (e.g. different languages and architectures). However, pinpointing whether a difference is cultural or not is difficult (e.g. is it culture, is it personal, is it environmental?). Culture is visible through comparison. Cultural influences are mostly recognized only when facing a different culture. There is no objective “culture” marker on things. Unlike temperature, a culture cannot be measured alone.

There is no consensually accepted definition of culture. Culture is a broad and fuzzy topic. Culture is easy to acknowledge (e.g. different languages and architectures), but hard to assert (e.g. is an observed difference cultural or personal?). Culture is not something that can be easily pinpointed: unlike emotions, individuals hardly recognize their own culture. As a concept, culture is difficult to define with precision.

In order to get a better idea about what culture is, we first introduce a few examples for grasping better what we discuss about. Then, we introduce the definition of culture that we rely on in this thesis. Next, we enrich this definition in introducing general aspects for characterizing culture. These aspects indicate the core features to expect and to look at when considering culture. Finally, given the number of aspects of culture, we clarify which ones are relevant with regard to this thesis and thus studied in the following chapters.

Examples of Cultural Manifestations

Cultural manifestations are relatively easy to observe when travelling. In reaching another country (the farther is generally the better), one can observe many unexpected things. Most of them are at least partly cultural. This section introduces basic examples of culturally-sensitive manifestations along the three core perspective for considering MASs: individual, collective and environmental.

The influence of culture on environments is the most visible. Culture influences how cities are built (e.g. is it a dense modern city? Are historical spots preserved?), architectures (e.g. Dutch houses made of bricks, French houses made of stone). One can see a lot of variation in decoration, in art (e.g. painting, sculptures), but also in symbols (e.g. visible flags, amount of road signs) and facilities (e.g. safe tracks for biking).

Culture influences individual behaviors. Superficially, one can easily see different habits: language, ways of eating and of greeting, walking speed, clothing. Looking into details, recurrent patterns can be seen: individuals tend to have different interests (e.g. focusing on money or on the family, on pleasure or on work, they enjoy different sports), they may care about different rules (e.g. compliance to the red light rule when driving, interrupting others when they talk or not, moving along together or not, bribing or not policemen).

Culture influences societies. Are people supporting the decisions of a dictator or requiring expressing their opinion through voting? Are there mechanisms for caring for poor or weak people or shall everyone care for him or herself? Are there syndicates for protecting subordinates from being abused? What about the rule system, are they tightly described or is it formed by informal rules of thumb? Is justice meant to punish forbidden behaviors or to bring equity back? Are policemen fine-givers or lesson-givers?

These are traditional cultural manifestations highlighted by literature. However, these manifestations can be partly discussed: are they cultural or caused by different environments and histories? In general, a bit of both. As a more "scientific" approach, one should consider situations where environmental context has little influence for observing better cultural differences.

Coffee rooms are relatively convenient² places for observing cultural differences. Again, many cultural manifestations can be observed. Just by indicating the country where the coffee room is placed in, theories can predict numerous expectable phenomena. Individuals can display different behaviors (e.g. quick coffee & back to work or time for a break and a chat? Which topics are discussed, weather, politics or sports?). Different culturally-driven social patterns can be observed (e.g. loners or groups? Are there rules for cleaning the cups?). Likewise, different cultural manifestations can be observed in the environment (e.g. are there seats and tables? What is the type of coffee machines? What are the quality and the speed of delivery? Is there a boiler for tea?).

Defining Culture

Culture is the key term of this thesis that we introduce here before extensively using it³.

²Because they are relatively unconstrained by the environment. As a side bonus, because they also serve coffee.

³As Voltaire used to say: "*Si vous souhaitez converser avec moi, définissez vos termes*", "*If you wish to converse with me, define your terms*".

Defining culture is a long-lasting challenge. As an illustration, Kroeber and Kluckhohn (1952) list several hundreds of definitions of culture. So far, no consensually accepted definition has been proposed. Multiple reasons can explain it. Culture is far from a solid object that can be sensed (e.g. like speed or temperature). Culture is a mental process, which is *hard to sharply predict or even measure*. Culture is only a collective tendency. It does not match strict rules of the form “if someone has culture *C*, then he or she will necessarily do *X*”. Even the concept of culture itself is *hard to ground*: unlike emotions individuals hardly recognize their own culture. *Culture is only observed through differences*, which implies noise caused by different contexts. Finally, culture is *complex*: culture can be related to indirect manifestations, but not with a few comprehensive sources. Even more complex, the consequences of culture indirectly impact on culture, leading to dynamic and cycling phenomena.

To that extent, in this section, we look for a definition that is simple, comprehensive and can be used to discriminate clear cases of the influence of culture. Hofstede et al. (2010a), the main theory of culture we rely on in this thesis, proposes: *culture is the collective programming of the mind that distinguishes the members of one group or category of people from others*. Culture is anchored within the minds of individuals. Culture is shared within groups or categories⁴ of individuals. Culture, through its influence, distinguishes the members of these groups from members of other groups. This definition is tightly related to a more comprehensive definition proposed by Kroeber and Kluckhohn (1952): *Culture consists of patterns, explicit and implicit, of behaviors acquired and transmitted by symbols, constituting the distinctive achievements of human groups, including their embodiment in artifacts; the essential core of culture consists of traditional (i.e. historically derived and selected) ideas and especially their attached values; culture systems may, on the one hand, be considered as products of action, on the other, as conditional elements of future action*. These definitions are enriched in later sections with more concrete aspects of culture.

General Properties of Culture

Culture can be related to numerous direct and indirect properties. This section briefly introduces these properties by relying on the MAS perspective: individual, collective and environmental perspectives.

Individual Perspective: Two main properties characterize culture at the individual level. First, culture is learned, and second, culture influences behaviors.

Culture is Learned: Fundamentally, culture is a form of learning. Culture results from the capability and the practical application of learning from other individuals. In considering very basic culture, individuals learn technical behaviors from each other (e.g. Studies investigating the root of cultures consider the *cyanistes caeruleus* species. This species learns how to open bottles of milk by looking at other birds doing it, highlighting a quickly spreading cultural phenomenon). Humans are capable of learning richer abstract mental patterns (e.g. ideas, songs, concepts) and can use more elaborated learning mechanisms (e.g. learning from books). In human societies, most of the cultural background is learned during the first years of life as indicated by Hofstede et al. (2010a). Afterwards, individuals can still learn new cultural “mental software” but the potential for change is more limited.

⁴Unrelated individuals with similar “mental programming” (e.g. individuals with similar training)

Culture Influences Individual Decisions: Culture influences other cognitive processes and indirectly individual decisions. Both influences are further studied in this thesis and are described in later sections. The influence of culture on cognitive mechanisms is further introduced in Section 2.2.3. The concrete manifestations of the influence of culture on individual decisions further detailed in Section 2.2.2.

The Influence of Culture Is Partial: Culture can relatively deeply influence individual decisions, but it does not *determine* individual decisions. Hofstede et al. (2010a) recognize two other factors as play: human nature and personality, as presented in Figure 2.1.

Human nature encompasses inborn decision capabilities that are shared by nearly all humans (e.g. reflexes, sense of danger, satisfying basic needs, basic logical thinking). Personality encompasses the influences on decisions that are specific to each individual (e.g. personal history).

Culture and Group Membership: Individuals tend to learn and express the culture of the group they belong to (e.g. rules, clothing, vocabulary, moral stands). Individuals tend to display cultural behaviors when they act in the context of this group (e.g. act according to the rules of the group).

Multi-Culturalism: Individuals can belong to multiple groups and learn multiple cultures. For instance, an individual can possess a culture for home (e.g. benevolent behavior), culture for work (e.g. obedient behavior) and a culture for sports (e.g. achievement-related behaviors). These cultures favor various behaviors that are expressed depending on the situation, particularly depending on the social context (at home, at work, at the sport club).

Culture and Expectations: Individuals expect their culture to be shared by other members of the group. Thus, individuals tend to expect that members of the group behave in accordance with the group culture (e.g. at the sport-club, there is a fair-play culture. Partners are expected to not cheat).

Collective Perspective: Culture, through its influence on individual decisions, impacts the constitution of social groups, societies and collective outcomes. Reciprocally, society

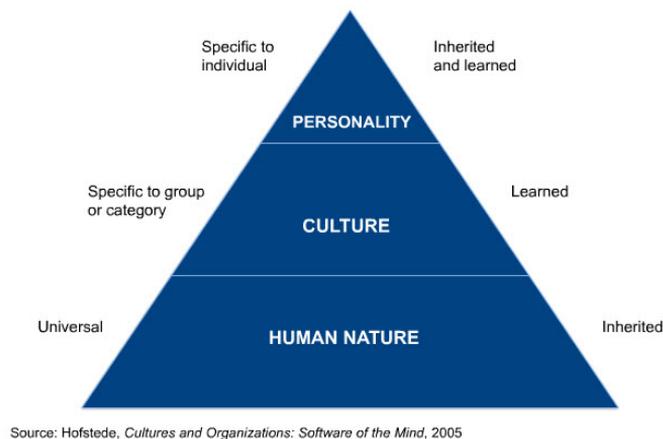


Figure 2.1 – Influence of culture on behavior, according to Hofstede et al. (2010a).

influences the type of culture that individuals can learn and how they can express their culture.

Culture Influences Interactions: Culture influences individual decisions and mutual expectations. Thus, it indirectly influences interactions (e.g. if fairness is culturally important, interactions tend to be less focused around checking for cheating. However, cheating in such a group tends to be more severely blamed). Furthermore, culture supports collectively shared mental structures (e.g. language, rules) that are crucial for successfully interacting.

Culture and Social Groups: Individuals sharing social groups tend to share the same culture. They tend to learn from each other, which homogenizes culture. Furthermore, social groups are generally related to mechanisms for enforcing cultural consistency within the members of the group (e.g. teaching newcomers, punishing individuals who do not comply with the culture of the group).

Reciprocally, culture can support the creation of social groups. Individuals sharing culture tend to interact more successfully with each other (e.g. shared language, rules) and thus to create social bonds.

Culture and Society: Culture and societies are tightly related. First, because culture supports the basic means for creating social bonds (e.g. common language, rules). Second, because culture supports collectively agreed-upon vision of the world. As indicated by Beck and Cowan (1996), this shared vision is crucial for building societies. For instance, if individuals share a culture where domination is at the core of interactions, then this culture is likely to support the occurrence of chiefdoms. Likewise, culture conveys technological knowledge, which indirectly dramatically impacts on the possible occurrence of societies (e.g. technology increases the food production, leading to an increase in the population size, influencing in turn the society). Reciprocally, as for social groups, society supports means for enforcing its cultural background (e.g. education, rituals and rules in line with culture).

Culture Indirectly Influences the Environment: Culture impacts the environment through its influence on individual decisions and societies. Three main visible influences of culture can be recognized. First, signs of group membership (e.g. clothing, architecture). Second, means for supporting cultural influences on individual decisions or societies (e.g. signs on the road, clocks on walls). Third, items aiming at carrying culture (e.g. books, typical souvenir items).

Mutual Influences: All previous aspects are interleaved with each other. The culture of an individual influences what this individual can learn next (e.g. rejecting ideas that are not in line with culture). Culture can influence behaviors that influence in turn the culture of other individuals (e.g. proselytism, teaching culture to others, blaming individuals for violating culturally-driven norms). Similarly, the influence of culture on individual decisions indirectly influences societies (e.g. social rules and structures are generally collectively complied with and this compliance is supported by culture). Reciprocally, societies can influence culture (e.g. educative and regulative structures). Likewise, influences of culture on the environment can change the situation of individuals (e.g. new production techniques) and thus the society and, in turn, culture (e.g. new environments require new individual perspective for handling them).

Core Decisions

The general background introduced in this section allows understanding better what culture is and to determine which aspects of culture are relevant in this thesis. Furthermore, we introduce here the main theory of culture we rely on throughout this thesis.

Sharpening our Focus: The aspects studied in this section indicate that culture is a complex topic. All aspects influence all aspects, including themselves. This difficulty is further increased by the lack of strong theoretical support for understanding in details these aspects and their mutual influences.

Not all these aspects are relevant with regard to this thesis. Since we aim at using culture in *practical* settings for supporting coordination, we should restrict the complexity of cultural models. In order to keep this complexity manageable, we focus on *a subset of the cultural aspects*. The remaining aspects should be sufficient for replicating (most of) properties of culture that are interesting with regard to our end-goal. Likewise, ignored cultural aspects should not impact (too much) on the properties of culture we want to reproduce.

In this thesis, we focus on replicating *the influence of culture on individual decisions and on collective outcomes*. In other words, we leave out cultural evolution. We assume a static culture provided by the system designer. Furthermore, while acknowledging its importance, we leave out the emergence of culturally-driven social structures for future research (as will be introduced in 8.3.3).

The set of studied cultural aspects is adequate with regard to the end-goal of this thesis. The influence of culture on coordination mostly results from the influence of culture on individual decisions (e.g. matching expectations, shared social paradigm) and indirectly on societies and collective outcomes. Aspects of culture that are left out are less critical for our goal: in human societies, cultural evolution takes many years, out of the scale of coordination. This feature can be interesting for building persistent systems, which is out of the scope of this thesis⁵.

Selected Theory: In this thesis, we focus on the theory of Hofstede et al. (2010a) as a core source of inspiration and validation.

Why focusing on a single theory? Numerous theories of culture are available and these theories are relatively fuzzy and abstract. To some extent, these theories even partly contradict each other. Thus, if we rely on many theories, we could easily “cheat” and use the theory and explanation we want for justifying any phenomenon we face⁶. In relying on a single theory, we possess a stable conceptual basis for building models, proposing explanations and matching models with expectations.

Why focusing on Hofstede et al. (2010a) and not another theory? This theory has numerous advantages. First, this theory is comprehensive, introducing a few core concepts instead of many sparse relationships, which is helpful for building models. Second, this

⁵As a side note, we explored this track at the start of this thesis, leading to Vanhée et al. (2013b,c). While being intellectually interesting, this direction is too premature. We were trying to replicate cultural dynamics, based on very limited theories, without having clear ideas about key mechanisms for replicating cultural behavior. This direction was hardly feasible, at least before having the current thesis.

⁶Here is an example from my personal experience about the negative consequences of multiplying theories. I ran a simulation for showing the first results of a simulation, still under development. Looking at the plots, I was told “Awesome! This simulation is perfectly in line with the theory A”. A few seconds later, I found out a bug, fixed it and obtained a completely different result. Here again, I was told “Awesome! This simulation is perfectly in line with theory B”.

theory is richly grounded with concrete visible manifestations of culture, which is helpful for validating models. Third, this theory shows relatively high robustness: this theory is (one of) the most empirically validated theories available and is related to numerous other studies.

2.2.2 Characterizing the Visible Manifestations of the Influence of Culture

The current section studies former research towards providing an initial answer to Research Question 1. This research question aims at capturing the visible manifestations of the influence of culture on coordination. To this extent, the current section studies the visible manifestations of culture in general. The former work depicting the influence of culture in situation of coordination is more specifically studied in Section 2.3 and conceptualized for being used within MASs in Chapter 3.

A first study of existing research reveals hundreds (if not thousands) of visible manifestations of the influence of culture. These manifestations are highlighted through many means such as case studies (e.g. comparing the same organization in three different countries as presented in D'Iribarne (1989)), polls, statistical comparisons (e.g. management by objectives works better in The United States than in France Hofstede et al. (2010a)) and examples (e.g. a French general very disappointed from being laughed at by subordinates).

In this form, former research depicting the manifestations of culture is impractical for building models. Thus, *in this thesis, we rely on general theories for explaining the influence of culture*. These general theories indicate abstract links for explaining (most of) the influence of culture instead of many low-level concrete relationships. Thus these abstract links offer a few general properties of the influence of culture that can be used more easily in this thesis.

Multiple reasons justify this decision. First, these general theories provide *comprehensive* explanations: they aim at explaining many links with a few concepts. Second, these theories are relatively *consistent*: while many studies can have different perspectives and slightly contradict each other, general theories limit these contradictions. Both aspects are crucial for building simple models, capturing most of the influence of culture, which is the aim of this thesis. As a drawback, these general theories offer slightly less accuracy: unlike studies, they do not go into details (e.g. very specific differences, sharp numbers). This drawback is relatively minor in our case: first because culture is too fuzzy to give too much care to details (culture indicates tendencies and details may be sensitive to other influences). Second, going into details generally implies increasing model complexity, which is something we want to avoid in this thesis.

*In particular, in this thesis we focus on Hofstede's Cultural Dimensions (HCDs) from Hofstede et al. (2010a)*⁷. Each HCD describes a generic abstract type of cultural influence (e.g. culture influences the importance given to power status). Then, this influence is conceptually and empirically related to many concrete manifestations (e.g. cultural importance given to power statuses is related to numerous manifestations such as higher obedience, leader-based organizational structures, communication forms for communicating with higher status individuals). In some sense, each HCD cluster is conceptually

⁷Other authors investigated similar cultural dimensions such as Hampden-Turner and Trompenaars (1993); Schwartz (1999)

and empirically correlated to cultural manifestations and explains these correlations with comprehensive concepts. Note that these correlations are statistical tendencies: not all cultural manifestations are visible in all matching cultures.

For sake of practicality, HCDs are measured by using a score, obtained through polls conducted at national level (e.g. in France, power statuses are culturally given a relatively higher importance than in The United States). These scores relate to expectable cultural manifestations. For instance, in countries where importance given to power is high (e.g. China, France), subordinates are more likely to comply with instructions without questioning the leaders than in countries where importance given to power is lower (e.g. The United States, The Netherlands). In this thesis, we often refer to extreme cultural scores in order to sharpen cultural differences. As a warning, these scores are statistical averages: all individuals from given country do not react the same to a similar situation.

Following subsections introduce in more details the HCDs presented in Hofstede et al. (2010a): Power distance (PDI-CD), Individualism (IDV-CD), Masculinity (MAS-CD), Uncertainty avoidance (UAI-CD), Long-term orientation (LTO-CD) and Indulgence versus Restraint (IVR-CD). For sake of disambiguation with Multi-Agent Systems, we add the suffix “-CD” after the standard acronyms for cultural dimensions. In this thesis, we leave out the IVR-CD, which is hardly connected to coordination (IVR-CD indicates the cultural influence on whether individuals are happy and spend time relaxing). These general descriptions are specified for the context of coordination in Chapter 3.

As a remark, in this thesis, we shorten the phrasings we use for referring to cultural dimensions. In our context, we want to regularly use assertions of the form: “In social groups/organizations where individuals share a culture that would achieve a high score on the PDI-CD, one can expect to observe that ...”. Instead, we use (much shorter) constructions of the form “In *high PDI-CD* organizations, ...”. We assume the latter phrasing to convey the meaning of the former.

PDI-CD: Power Distance

PDI-CD indicates the cultural importance given to power statuses. PDI-CD can be seen as the cultural response to the following social dilemma: should leaders be the ones leading the group or should everyone participate to collective decisions? Both answers make sense but conflict with each other. Culture influences the relative preference to one or the other and on indirect individual and social consequences. The same remark applies for following HCDs.

In high PDI-CD societies, individuals tend to value modesty, obedience and discipline. Empowered deciders should make decisions for the group and others should comply with these decisions. Both roles are strongly dependent on each other. This culture tends to give a lot of importance to decision makers supporting strong social inequalities.

In low PDI-CD societies, individuals tend to value equality, autonomy and self-direction. Power relationships between individuals tend to be weaker. Individuals tend to prefer democratic decisions (decision maker should make decisions in line with opinions of others). As a consequence, all tend to be responsible for decisions and their consequences. Furthermore, people tend to be more independent from decision-maker (e.g. more initiative).

E.g.: China, Russia (high PDI-CD) opposed to Scandinavian countries (low PDI-CD).

IDV-CD: Individualism versus Collectivism

At the core, IDV-CD can be related to the relative importance between the satisfaction of individuals and of groups (e.g. shall individuals serve groups or shall groups serve individuals? Hampden-Turner and Trompenaars (1993)). This relative importance is particularly visible in the commitment of individuals to groups (e.g. strong relationships with group members, compliance with group identity and rules).

High IDV-CD societies tend to value freedom and independence between individuals. Individuals tend to focus on their own interest and expect others to do the same. When acting in groups, individuals tend to gather by common interest. They tend to keep their own identity, with limited display of group identity (e.g. symbols, rules). Individuals tend to create limited ties with others.

In low IDV-CD societies, individuals tend to create strong mutual dependences with the social groups they belong to. Groups protect individuals and individuals should protect the group and other individuals of the group. The underlying idea is that everyone is more satisfied when they all work in the same direction. These societies tend to support high compliance to group identity (e.g. symbols, goals and rules, defending the group interest). Furthermore, individuals tend to create strong emotional and social ties with each other, further reinforcing this integration within groups.

E.g.: Anglo-Saxon countries (high IDV-CD) opposed to South American and Asian countries (low IDV-CD).

MAS-CD: Masculinity versus Femininity

MAS-CD can be related to the cultural answer to the following dilemmas: “assertiveness versus modesty” and “mastery versus harmony” and “equity versus equality”.

High MAS-CD societies tend to support a preference for the first answer of these dilemmas. In high MAS-CD societies, individuals tend to value individual performance. Individuals should aim at being the best possible (e.g. success, excellence, performance, competence). Good performers should be rewarded, equity is important (e.g. rewards based on commission or tips). Individuals are culturally driven to master their environment. This mastery is socially displayed through assertiveness, showing mark of success. Indirectly, this race for performance tends to support high competition.

Low MAS-CD societies are characterized by a preference towards avoiding conflicts. Individuals should avoid opposing each other, preferably working together. They tend to be more willing to cooperate, make compromises, look for consensus, support and care about each other. In social interactions, individuals tend to be modest. They tend to avoid raising conflicts by not matching their promises. They tend to prefer remaining modest in order to avoid disappointment. Similarly, such societies support equality as a basis for spreading rewards for avoiding conflicts (e.g. competition, jealousy).

E.g.: Scandinavian countries (low MAS-CD) versus Italy, Japan (high MAS-CD).

UAI-CD: Uncertainty Avoidance

UAI-CD can be seen as the influence of culture on the sensitivity towards anxiety raised by uncertainty.

In high UAI-CD societies, individuals tend to be more sensitive to anxiety. They tend to take more action in order to lower it. In general, individuals tend to rely on standardization when making decisions. Standardization tends to offer a relative certainty and stability about consequences of decisions. Furthermore, standardization is socially desirable because it also provides a relative certainty to others. When communicating, individuals tend to rely on standards and to communicate more for handling unexpected situations. Finally, experts tend to be given a relatively high importance. From a general perspective, individuals tend to prefer keeping low uncertainty and expect from others that they also keep low uncertainty. High reliance on standardization generally leads to strong assumptions about the environment (e.g. incoming products are expected to be normal).

In low UAI-CD societies, individuals tend to be less sensitive to uncertainty; individuals tend to require more uncertainty before reacting. Individuals are more likely to consider each problem independently and look for *ad-hoc* adaptive and creative resolutions. They tend to avoid creating assumptions about the state of the environment and avoid making sure of the actions of others through tight planning. In some sense, they locally re-create the certainty they need for achieving their goals.

E.g.: Russia, Germany, Japan (high UAI-CD) versus Switzerland, Sweden, Denmark (low UAI-CD).

LTO-CD: Long-Term Orientation, Normative versus Pragmatic

LTO-CD can be tightly related to the relative importance given to short-term versus long-term consequences of decisions. In high LTO-CD societies, long-term consequences tend to be given more importance than immediate ones and vice-versa for low LTO-CD societies.

High LTO-CD societies tend to give high emphasis to long-term consequences. Individuals tend to care more about learning, maintaining good health, thrift, building and maintaining strong social networks. Long-term pragmatic learning tends to be preferred. Individuals tend to accept failures more easily and consider them as a learning experience, praising perseverance.

Conversely, in low LTO-CD societies individuals tend to focus on more immediate rewards. Individuals tend to be more ready to spend money and build relatively limited relationships. Individuals expect their situation to change often. They are more likely to develop mechanisms for swiftly learning, apply and exchange simple rules for making decisions. They tend to rely on logical thinking which quickly provide relatively good outcomes with limited investment for learning. Failures are seen as a “loss of face” and are strongly avoided and hidden.

E.g.: China, India (high LTO-CD) versus Canada, Great Britain (low LTO-CD).

2.2.3 Core Mechanisms of Culture

In this section, we investigate former research for providing an initial answer to Research Question 2.a. This question aims at determining the core mechanisms of culture from which (most of) concrete manifestations of the influence of culture can be derived.

Hofstede et al. (2010a) introduce two central mechanisms influenced by culture that can explain most of cultural manifestations: *values* and *practices*. These mechanisms are

introduced in the “onion diagram”, Figure 2.2.

As a warning, this connection is only partial: the relation between culturally-sensitive mechanisms and cultural manifestations, while being acknowledged, remains informal. Current research does not describe with precision culturally-sensitive cognitive mechanisms, how culture influences these mechanisms and how these mechanisms lead to concrete behaviors. Additional research is required for answering Research Question 2.b. As a note for Section 2.2.4, the lack of formal connections hints why theories and subsequent models rely on empirical correlations (e.g. HCD style) and not on cognitive mechanisms for explaining or replicating cultural manifestations.

Following subsections briefly introduce each cognitive mechanism, its relation to both culture and decisions.

Practices

Practices constitute the most explicit, visible and concrete influences of culture on decisions. Hofstede et al. (2010a) describe three types of practices that are particularly sensitive to culture: *symbols*, *heroes*, *rituals*.

Symbols correspond to shared symbolic concepts or items and their interpretation (e.g. technical terms, languages, flags, clothing). Heroes are people that are known to be of great importance and should be mimicked (e.g. national heroes). Rituals are behaviors that are recognized to be important for social recognition (e.g. how to greet, ceremonies). Additional culturally-sensitive practices can be considered (e.g. rules, roles), but tend to overlap with non-cultural influences.

Culture influences the range of practices that are acquired and applied by an individual. Individuals sharing a culture are expected to be aware and to some extent comply with these practices (e.g. citizens from a country are expected to know the color of their flag, individuals practicing the same sport are expected to know the rules).

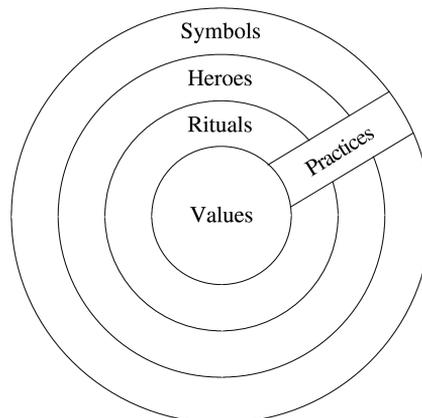


Figure 2.2 – Onion diagram presenting the internal structure of cultures according to Hofstede et al. (2010a)

Values

Values (Miceli and Castelfranchi (1989)) are a rich and complex cognitive mechanism. Briefly, values indicate what is “good”, or “good” features of things. This indication is relatively accurate with regard to one’s welfare (e.g. lying generally leads to undesirable outcomes), but remains irrational (e.g. telling the truth because “honesty is good”, but no one benefits from this). In practical settings, values are used for fastening decisions: since they indicate (generally) good features of decisions, values can be used as “heuristics” (e.g. quickly discarding the decision to lie before considering all possible consequences of lying). A more complete description will be proposed in Section 4.1.2.

Values (e.g. honesty, creativity, cleanliness, efficiency) have a more subtle influence than practices. Values influence a very wide range of decisions but they are much more implicit and difficult to pinpoint (e.g. one can say “this is the ritual” but hardly “this is the value”).

In considering their relation with culture, culture influences the type of values that individuals tend to possess. Furthermore, culture influences the relative importance between values. This importance matters when multiple values propose evaluations that enter in conflict (e.g. “report the fault of my boss” if “good” according to the value of honesty and “bad” according to the value of obedience. The relative importance given to these two values influences the final decision).

Global Properties of the Influence of Culture on Cognitive Mechanisms

The influence of culture on cognitive aspects as depicted by the onion diagram has multiple global properties. First, internal layers (influence of culture on values) tend to convey more influence on decisions. Internal layers tend to influence a wider range of decision situations than external layers. Internal layers are harder to learn and to change. Internal layers are more fuzzy and harder to pinpoint than external ones. Finally, the influence of culture on cognitive mechanism tends to be internally and globally coherent. For instance, on the same level, culture tends to support both values of “obedience” and “respect social order” than “obedience” and “autonomy”. On different levels, culture favoring the value of “obedience” is more likely to support symbolic forms of politeness for power (e.g. “tu” and “vous” in French). In some sense, culture tends to form unified influences over practices and values.

2.2.4 Computational Models of Culture

In order to use culture for supporting coordination in MASs, we need to be confident that the influence of culture can be replicated within MASs. In order to achieve this goal, we aim at understanding how to build models of culture. This direction is investigated through Research Question 2 and more specifically through Research Question 2.b.

An initial research indicates that several models of culture are available. In this direction, two main categories of models can be considered. First, models of culturally-sensitive mechanisms, as identified in Section 2.2.3. Second, models aiming at replicating cultural manifestations as identified in Section 2.2.2. The intersection between both categories is limited, more details about why are provided in the summary.

Each following subsection introduces either models of culturally-sensitive mechanisms or models for replicating cultural manifestations.

Modeling Culturally-Sensitive Cognitive Mechanisms

The current section introduces available computational models of cognitive mechanisms which are known to be culturally-sensitive, as introduced in Section 2.2.3. To our knowledge, these models are not connected to the influence of culture and are not used for replicating manifestations of culture.

Practices: practices can be related to numerous available frameworks. In particular, these frameworks can be used for coordinating MASs. Symbols can be represented by data structures, ontologies (Gruber (2005)) or communication languages (Wooldridge (2002)). Rituals can be related to numerous coordination mechanisms, particularly norms (Boella et al. (2006)) or protocols (Wooldridge (2002)). Heroes have received less attention but they can be modelled by inheritance mechanisms between agent decision processes.

Values: Values have also received attention. Nevertheless, models of values remain partial because they are too implicit and are tightly embedded within decisions, as depicted by Miceli and Castelfranchi (1989).

Available models of values (such as the one proposed in van der Weide (2011)) aim at capturing several core properties of the influence of values on decisions. In particular, this model can replicate one of the core aspects of values: when facing alternatives that can hardly be rationally sorted, individuals are driven towards making value-compliant decisions. In this model, the value-systems are represented by a preference order (Brafman and Domshlak (2013)) over outcomes of decisions. Further details about this model are provided in Section 4.2.1. In a similar direction, Antunes and Coelho (1999) propose an informal model of values for evaluating goals and plans in order to determine which action to perform.

Modeling Cultural Manifestations

Several former models aim at replicating concrete manifestations of culture within MASs. All available models aim at building simulations: they define a set of decision situations and aim at observing behaviors that match studied cultural properties.

Two main categories of models can be considered. The first category focuses on replicating the dynamics of the influence of culture. The second category focuses on replicating the influence of culture on individual behaviors and societies. Each category of models tends to focus on limited aspects and to blur the others (e.g. models replicating cultural evolution tend to have limited models of individual behaviors and societies).

Modeling Cultural Dynamics: This category of cultural models of culture aims at replicating properties of cultural evolution and cultural propagation (e.g. adoption of technology, inter-cultural conflicts, cultural polarization). These models are out of our focus and are only briefly introduced here for sake of completeness.

At the core, these models rely on a similar representation of culture as presented by Axelrod (1997); Epstein and Axtell (1996). Each agent possesses a set of cultural mental patterns, represented by a “cultural vector” (e.g. memes). This cultural vector influences in turn behaviors of agents (e.g. teaching/learning bits of culture to/from others, making decisions not directly related to cultural propagation such as fighting, growing food).

Modeling Cultural Influences: This category of models aims at accurately replicating the manifestations influence of culture, assuming a static culture. In particular, these models generally focus on either replicating individual decisions or collective outcomes.

In most of the cases, culture is modeled by integer variables representing HCD scores. These variables are integrated within numerical formula which are used for making decisions (e.g. need to evaluate the status of another individual, based on a sum of parameters that includes amongst others “ $power_diff \times PDI - CD$ ”, where *power_diff* measures the difference in terms of power between the two individuals). Nevertheless, in certain cases, such as Dechesne et al. (2012); Mc Breen et al. (2011), culture is modeled by the selection of more specific decision rules (actually, in implicitly using values within the system). The link between these rules and underlying cultures implicit within the model and is justified out of the model.

A first subcategory of models aims at replicating detailed influence of culture on *individual* decisions. These models are mainly used for enacting culturally credible virtual characters, such as the work of Degens et al. (2013); Endrass et al. (2013). In general, these models consider relatively rich decision models (e.g. emotions, status, power, trust). They also tend to be related to a rich set of individual actions (e.g. speech acts, gestures, clothing).

A second subcategory of models focuses on replicating collective phenomena that result from the influence of culture on individual decisions. In this approach, individual decisions are kept relatively simple, towards finding the minimal mechanisms for explaining emergent phenomena (as indicated by Epstein (1999)), for keeping computational complexity low and for avoiding noise. In this direction, Hofstede et al. (2010b) use HCDs for influencing individual decisions and observing collective phenomena observed in trading organizations (e.g. organizations sharing the same culture tend to achieve better outcomes, high LTO-CD organizations tend to be more forgiving, leading to longer-term partnerships).

We contributed to both tracks in multiple extent. We proposed models of the influence of culture on decisions and coordination in Vanhée et al. (2014a) as will be presented in Chapter 5. We explored models of the influence of culture on trust and trust building for multiple application domains (robotics, virtual agents) in Borit et al. (2013, 2014a,b). We contributed to this track from an agent-based perspective in the playground model, presented in Hofstede et al. (2014). In this model, agents represent children. Culture influences their social behavior (e.g. agents tend to be more or less rough with each other depending on MAS-CD). This influence leads to different emerging social patterns, in accordance with observations from human playgrounds (e.g. in high MAS-CD societies, genders tend to more segregated).

As a side note, relatively rich and interactive models of culture are proposed by the game industry for modeling AI players. In our opinion, one of the best available models of culture is proposed by Civilization V. In this game, each AI civilization is related to a set of cultural traits (e.g. forgiveness for past wars, aggressiveness, willingness to trade, reciprocity). These traits then influence decisions made by the AI in general (e.g. focusing on gathering luxury or military resources, city development) and particularly in the context of diplomacy (e.g. tendency for retaliating, for preparing surprise strikes against allies, for being generally peaceful).

Strong Boundaries

This section introduces existing research about computational models (1) for replicating cultural manifestations and (2) of core cognitive mechanisms that are sensitive to culture.

As a surprise, both categories of models does not overlap. This can be explained by considering available theories and the purpose of these models. Models of type (1) are simulation models: they aim at replicating credible cultural behaviors for limited range of decisions. Thus, they can rely on HCDs that indicate credible cultural expectations. Building simulations do not require type (2) models: these models can handle more decisions (useless because decisions are bounded) and require to design a more complex machinery. The other way around, models of type (2) are used either for supporting complex individual decisions (in the case of values) or for supporting coordination (for practices) but nor for replicating cultural manifestations.

In this thesis, we cross this limit by *combining both categories of models*. On the one hand, our model should replicate manifestations of culture. On the other hand, we want our decision model to be rich enough, because agents have to make complex decisions for handling a wide range of coordination situations. Thus, we need to combine both types of models.

Our model of culturally-sensitive decision aspects is proposed in Chapter 4. We show that this model replicates cultural behaviors in Chapter 5 and Chapter 6. This model is then used as an evidence for showing that the influence of culture can be replicated in MASs and handled for supporting coordination in Chapter 7.

Conclusion

This section investigated former research for introducing a general background (definitions, concepts, properties) for understanding better culture. Based on this background, we indicated which aspects of culture we focus on in this thesis (replicating the influence of culture on individual decisions and collective outcomes) and which aspects we leave out (learning, evolution of culture). These aspects allow limiting the complexity of the models we need to produce and creating stronger expectations about how to use culture for supporting coordination.

Then, we investigated former research for conceptualizing better the manifestations of culture. In this direction, we introduced Hofstede's Cultural Dimensions from Hofstede et al. (2010a). HCDs indicate comprehensive categories of conceptually and empirically correlated cultural manifestations. We also introduced the core mechanisms of culture. The influence of culture is particularly visible through two cognitive mechanisms: values and practices. Finally, we introduced computational models of culture and of culturally-sensitive cognitive mechanisms.

This study raises core limitations about the current state of our understanding about what culture is. At the moment, we cannot tell exactly what culture is. Culture is informal, fuzzy, intangible, implicit and only visible through (statistical) comparison. Two core mechanisms of the influence of culture are recognized (values and practices), as well as some categories of cultural manifestations. Nevertheless, *we cannot rely on exact links* such as "if $PDI - CD > 50$ then subordinates always agree with leaders". We can only rely on relatively abstract general and relative links "in high PDI-CD, subordinates tend to agree

more often with leaders”. For sake of comparison, we are nowhere near very detailed theories such as Newtonian physics, with sharply detailed relationships and which can be easily verified.

These limitations do not prevent achieving the goal set by this thesis. Indeed, we do not need to replicate sharply all cultural manifestations. First, because not all cultural manifestations are useful with regard to our goal (e.g. replicating the influence of culture on greeting by bowing or shaking hands has little interest for supporting coordination). Second, because aiming at reproducing in details all cultural manifestations would increase model complexity. To that extent, we do not need to understand and model culture up to the finest details. Thus, *in this thesis, we focus on replicating the core aspects of and manifestations of the influence of culture.*

Pushing the Limits

In this thesis, we enrich the theories proposed in this section. In Chapter 6, we will expand current theories explaining cultural manifestations presented in Section 2.2.2. In particular, we propose a generic explanation for understanding how the influence of culture on individual decisions leads to the influence of culture on collective outcomes.

In Chapter 4, we will expand the models of culture presented in Section 2.2.4. More specifically, we will propose an architecture that can both build models capable of handling a wide range of decisions (in particular, complex decisions in situation of coordination) and replicating cultural manifestations. For sake of illustration, two simulations based on this model will be proposed in Chapter 5 and Chapter 6.

2.3 Relating Culture & Coordination

Previous sections introduced independently the necessary background for better understanding the two core topics of this thesis: culture and coordination. The current section studies in details the intersection of both topics.

As for previous sections, we first introduce a general background for better understanding the main relationships between culture and coordination. Then, we study more in details available research towards answering more specifically our research questions. For answering Research Question 1, we study available research depicting the manifestations of the influence of culture in situations of coordination. For answering Research Question 2, we study available research introducing the core mechanisms of the influence of culture at play in situations of coordination and available computational models. In order to answer Research Question 3, we study existing research that investigates how to use culture as a tool for supporting coordination.

The general background is introduced in Section 2.3.1. Manifestations of culture in situation of coordination are introduced in Section 2.3.2. The key mechanisms of culture in situation of coordination are introduced in Section 2.3.3. Methods using culture for supporting coordination are introduced in Section 2.3.4.

2.3.1 General Background: Culture & Coordination

The current section specifies the background for understanding culture presented in Section 2.2.1, from the general context to coordination context. The structure is the same as in Section 2.2.1: we first introduce definitions and aspects of culture in the context of coordination and then we introduce which aspects we focus on in this thesis.

Specifying Definitions and Aspects of Culture, in the Context of Coordination

Culture in the context of coordination share many aspects in common with culture in general. The definition is the same: culture is still mental influences that differentiates groups or categories of individuals. Similar key aspects can be found. From an individual perspective, culture is learned, influence partially individual decisions, is used for determining group membership, individuals can possess multiple cultures and culture influences expectations. From a collective perspective, culture influences interactions, social groups and societies. Finally, culture indirectly influences the environment.

In addition to these general aspects, more specific aspects are pushed forward in the context of coordination.

Levels of Culture: In the context of coordination three main types of culture, corresponding to three “levels” can be found. First, a *general cultural background* (e.g. national culture detailed in Section 2.2.2). This background is shared by all members of the group and out of the group. An *organizational culture*, shared by members of the organization. This background is mostly learned within the organization and related to organizational activities. A *local culture*, shared by members of similar “groups” (e.g. working units, groups such as accountants vs. production).

These various types of culture influence decisions slightly differently. The general cultural background level conveys a general influence on decisions. This form of culture particularly influences values, which are in turn expressed within the organization. The last two are more driven by activities and processes of the organization and units of the organization. These latter types of culture tend to mostly influence practices.

The Influence of Culture is Partial: Culture mostly influences *informal* (e.g. rarely takes over formal constraints such as direct orders, rules) and subjective decisions (e.g. not influencing obvious decisions). As indicated by Hofstede et al. (2010a): “The less an activity is determined by technical necessity, the more it is ruled by values and thus influenced by cultural differences.”

Culture Influences Individual Decisions: In practice, restrictions caused by “technical necessity” in the context of coordination are minor (particularly when aiming at achieving flexible coordination as indicated in Section 2.1.2): culture has a major influence on decisions.

Hofstede et al. (2010a) indicate numerous examples of the influence of culture on decisions in the situation of coordination (e.g. who should make this decision, which goal to adopt, how tight plans should be, how to handle this issue?). These examples are numerous and can hardly be introduced exhaustively. They are related to each other and general cultural influences in Section 2.3.2.

Culture and Interactions: Culture indirectly influences interactions. Again, Hofstede et al. (2010a) provide numerous examples: who manages task allocation and workload leaders or subordinates? Are individuals focusing on their own performance or be willing

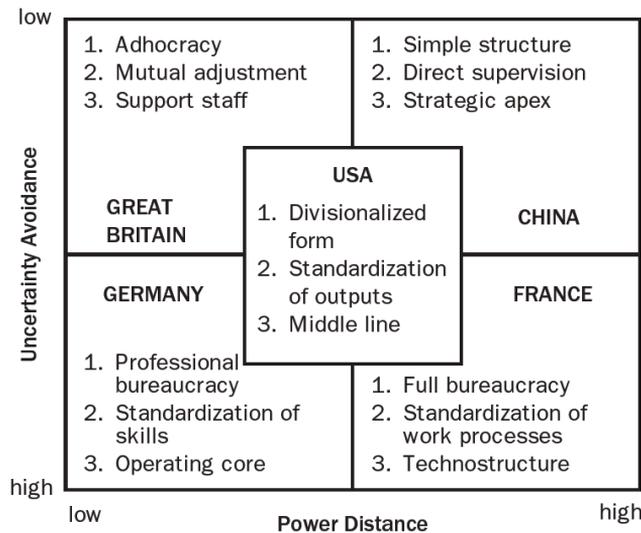


Figure 2.3 – Preferred Mintzberg canonical organization depending on culture according to Hofstede et al. (2010a)

to support failing others? How are conflicts solved? Changing rules, letting leaders settle the conflict, informally training conflicting individuals to work better together?

The influence of culture on interactions is partly “conscious”. Individuals are explicitly culturally-driven towards preferring a form of coordination (e.g. relying on a leader is better). However, they are also implicitly culturally-driven towards the occurrence of this interaction pattern (e.g. complying with decisions of leaders, without assuming that leadership is better).

Culture and Formal Structure: Culture drives the occurrence of interaction patterns that can be supported or not by the formal coordination structure. Individuals tend to prefer a match between the two. Furthermore, formal structures that match a given culture tend to occur more frequently in practice (e.g. in a culture promoting high UAI-CD, coordinating with rules is more frequent). This can be explained because changing organizational structure generally involves informal influences.

In this direction, Hofstede et al. (2010a) comprehensively illustrates in Figure 2.3 the key relationships between cultural background and preferred canonical structure from Mintzberg (introduced in Section 2.1.2). This figure illustrates the type of coordination pattern that tends to arise and to be supported depending on culture. In particular, reading this figure along the UAI-CD axis indicates the influence of culture on the preferred degree of standardization (full bureaucracy vs. simple structure; professional bureaucracy vs. adhocracy). Reading the figure along the PDI-CD axis indicates the influence of culture on the preferred degree of centralization around leaders (simple structure vs. adhocracy; full bureaucracy vs. professional bureaucracy).

This figure is at the center of this thesis. This figure captures the core cultural manifestations that we will simulate in Chapter 5 and Chapter 6. Keeping this figure in mind helps understanding the core manifestations of culture we will replicate in these chapters and

for better understanding the mechanisms that lead to our results.

Framing Our Exploration

As indicated in Section 2.2.1, we restrict the aspects of culture we aim at modeling in this thesis. As indicated in Section 2.2.1, this thesis leaves out aspects of learning a new culture and thus cultural evolution; culture is assumed to be static and provided by system designers (or coordinators).

In this section, we introduced new aspects of culture that are visible in the context of coordination. We aim at capturing all these aspects at the exception the aspect titled “levels of culture”. *In this thesis, we only consider the general cultural background.* We do not investigate organizational and local cultures.

Multiple reasons justify this decision. Fundamentally, the general cultural background provides the highest flexibility. This background influences values that in turn support a wide range of decisions. In comparison, organizational and local types of culture are much more superficial. They are mostly practices, highly driven by existing coordination methods but less appropriate for generating adaptive coordination. Furthermore, from a more practical perspective, general cultural backgrounds have received much more attention until now, both from conceptual and empirical perspectives (e.g. extensive descriptions from HCDs).

2.3.2 Manifestations of the Influence of Culture on Coordination

Research Question 1 investigates the manifestations of the influence of culture on coordination.

As introduced in Section 2.2.2 for culture in general, many manifestations of the influence of culture on coordination can be captured by HCDs. HCDs indicate comprehensive categories of empirically-related concrete manifestations of the influence of culture in general. In addition to this general influence, each HCD can be related to concrete cultural manifestations in the context of coordination. Many relations between HCDs and these manifestations can be proposed. Given that we will conceptualize these manifestations in the context of MASs in Chapter 3, we leave their detailed introduction in the current section.

In addition to these concrete manifestations, Hofstede et al. (2010a) indicate general patterns of the influence of culture on coordination. In particular, Hofstede et al. (2010a) show the presence of strong correlation between culture, individual decisions, interaction patterns, formal organizations and organizational performance.

As an example of such correlation, consider the case of a culture promoting high PDI-CD manifestations. In this culture, (1) individuals tend to prefer relying on leaders when making decisions; (2) interactions are more likely to be focused on leader-subordinate relationships; (3) formal organizations tend to possess leaders; (4) organizations are more likely to achieve higher efficiency in simple environments, lower efficiency in complex environments and lack robustness in comparison with a culture promoting low PDI-CD manifestations.

As a further illustration, in this culture, trying to coordinate without a powerful leader (e.g. management by objective requiring a fair bargain between leaders and subordinates)

tends to be both negatively accepted by individuals, be poorly applied (e.g. no decisions are made or this method is discarded for relying on standard leader-subordinate relationships), leading to poor performance (e.g. lower performance in comparison with the same structure in low PDI-CD organizations). As Hofstede et al. (2010a) indicated: “Assuming that management ideas are universal is naive”.

As a final notice, available theories do not explore Research Question 1.c (explaining the influence of culture on the collective level based on the influence of culture on the individual level). Available theories focus on the influence of culture on either the individual or collective level, without trying to explain in detail the relationships between these two levels. The answers to Research Question 1.a and Research Question 1.b provide an initial answer for Research Question 1.c, but no clear answer is available, yet. We provide an answer to this question in Chapter 6. The need for explaining this phenomena ourselves is an example that illustrates why this thesis is exploratory.

2.3.3 Key Mechanisms of Culture Involved in Situation of Coordination

Research Question 2.a aims at capturing the key mechanisms of the influence of culture in situation of coordination and Research Question 2.b aims at modeling these mechanisms.

General mechanisms of the influence of culture introduced in Section 2.2.3 also apply in the context of coordination. Thus, in the context of coordination, culture influences the symbols, heroes, rituals and values of individuals, which influence in turn their decisions. About modeling these mechanisms, there is nothing to add more than what we introduced in Section 2.2.4.

2.3.4 Using Culture as a Means for Supporting Coordination

Former research provides an initial answer for Research Question 3. Former research (only) explores how to use culture for coordinating in human societies, which is the core source of inspiration of this thesis. These studies indicate two core properties.

Multi-Cultural Interactions are difficult to handle. These interactions involve individuals with different cultures and thus with different conceptions about what is important (values) and how to interact with others (practices). Thus, such interactions generally lead to failures. Two main solutions are proposed: including cultural negotiators (individuals capable of behaving according to, understanding and managing both cultures at the same time) or relying on mono-cultural crews. Given the difficulty of managing multi-cultural interactions, *in this thesis, when aiming at using culture for coordinating agents*⁸, *we assume that all share the same culture*. We aim first at capturing core features of culture before investigating more specific and difficult features of culture. Nevertheless, we acknowledge that multi-cultural crews can be interesting for optimizing performance further (e.g. having one culture per working unit. This culture fits the coordination techniques which fits the local problem to handle), as will be introduced in Section 8.3.3.

⁸In this thesis, we do not always use culture for coordinating in all chapters. In particular, Chapter 6 simulates the influence of culture in human-like societies, where culture is not used as a means for coordinating agents. In particular, this chapter illustrates the negative consequences of multi-cultural interactions. These negative consequences further justify the current limitation!

As a note, issues raised by multi-culturalism partly explain limitations of current methods for supporting flexible coordination in MASs. Current methods do not take into consideration the “culture” (or a synonym) of their agents. To that extent, multi-culturalism and its issues can implicitly occur within MASs, leading to coordination failures that are hard to explain.

As a second observation, theories (Hofstede et al. (2010a); Steffan (1997) indicate *the importance of the match between culture, coordination mechanism and other coordination aspects*. As introduced in Section 2.1.3, coordination methods need to be attuned to the coordination problem (e.g. highly ruled systems are adequate in simple environments for achieving high flexibility) and as introduced in Section 2.3.2, the coordination method also needs to be attuned to match culture (e.g. better interact when a leader makes collective decisions in a high PDI-CD culture). In conclusion, all aspects need to match in order to achieve performance. Thus, when using culture for supporting coordination in MASs, coordinators should make sure that culture, coordination methods and coordination problems match each other. For instance, in crossing Figure 2.3 and canonical organizations from 2.1.2, the best solution for handling simple and dynamic environments consists in using a simple structure in a high PDI-CD and high UAI-CD culture.

In conclusion, human culture can be used as a tool for supporting flexible coordination. Culture indicates commonly-shared (and strongly expected) abstract solutions for “working together” by defining what “harmony” means. As introduced in Section 2.1.1, “working together” and “harmony” are at the core of the definition of coordination. Thus, culture, by clarifying what “harmony” means is crucial for flexibly establishing coordination. One step further, since culture is shared, individuals sharing the same culture share a common idea of what “harmony” means. Thus, these individuals agree on what is important when coordinating (e.g. following rules; orders) and common coordination patterns (e.g. leaders decide for subordinates) but also about adequate coordination-related decisions at the individual level (e.g. comply with instructions from leaders) and expectations about coordination-related decisions from other individuals (e.g. instructions from leaders are expected to be complied with). Thus, if their culture matches the given problem, they have all the cards at hand for adequately generating coordination and acting according to coordination once it is set, allowing to handle flexibly a wide range of problems.

The rest of this thesis is dedicated to replicating this human phenomenon within artificial societies in order to use culture as a means for supporting flexible coordination.

2.4 Chapter Summary

This chapter introduced former research about the two core topics of this thesis: coordination and culture, first independently and then their intersection. Each time, we introduced a general overview of the topic and we focused on how this former research can help answering our research questions.

As a brief summary, coordination consists, for a coordinator, in “harmoniously” solving dependencies between agents in order to achieve a given goal. Coordination is a complex task that requires taking numerous aspects in consideration (e.g. environment, performance, degree of autonomy). In our case, we particularly aim at achieving flexible coordination, where a given coordination can handle a wide range of possibly evolving goals

and environments. Through a review of existing methods for achieving flexible coordination, we discovered that most appropriate methods rely on the system (generally, through agents) for dynamically generating coordination from within the system, when facing a concrete situation. We also discovered that with these methods agents miss to a core aspect for achieving flexible coordination: these methods do not describe clearly what is the “harmony” that agents should pursue when being coordinated.

Culture is actually a good candidate for defining harmony. Fundamentally, culture is a mental influence that distinguishes groups from each other. Culture deeply influences the way individuals interpret and react to the world around them, particularly when considering social interactions. Culture is shared, providing a basis for establishing adequate mutual understanding and expectations.

Culture impacts coordination in human societies. Coordinating individuals with different cultures is particularly difficult: their different perspectives of the world lead to misunderstandings, misalignment of expectations, distrust and thus to failures. Likewise, culture and techniques for coordinating individuals need to match. Otherwise, coordination techniques in place are misapplied or discarded, leading, again to failures. Thus, culture is an important factor for successfully coordinating: in considering them properly, culture can be used as a tool for supporting coordination.

The next step consists in replicating this definition of harmony within MASs. In this direction, we identified two key mechanisms of culture: values and practices. Then, we explored available solutions for modeling them. Several models are available, but additional work remains to be done in this thesis in order to adequately replicate the influence of culture on them and use them for supporting coordination in artificial societies.

Regarding following chapters, the study of former research conducted in this chapter indicates partial answers for our research questions and what remains to be investigated in following chapters, framing our contribution. Available research partly covers Research Question 1.a and Research Question 1.b, by describing the influence of culture on individuals and societies in general. Nevertheless, these descriptions are too general for our goal: we want more details about the influence of culture in the specific case of coordination in order to accurately replicate this phenomenon. In this direction, we will further specify this link in Chapter 3. Regarding Research Question 1.c, this question has not been covered in details by former research. This question will be further explored through simulations in Chapter 6. Regarding Research Question 2.a, former research acknowledge some key mechanisms of the influence of culture (practices and values). Nevertheless, the influence of these mechanisms on the occurrence of flexible coordination needs to be further explored. Regarding Research Question 2.b, former research does not explore how to model culture in order to support flexible coordination. These mechanisms will be explored and modeled in Chapter 4 and will be tested against expectations from social sciences in Chapter 5. Finally, former research does not study Research Question 3. This question will be studied in Chapter 7.

First Answer for Research Question 1: *How Does Culture Influence Coordination in Human Societies?*

Culture influences many aspects of individual decisions (e.g. rely on leadership, rely on rules, aim for long-term goals). This influence leads in turn to many manifestations of cul-

ture on collective outcomes (e.g. culture influences whether organizations achieve good performance in simple environments).

In order to handle this complex influence more practically, theories such as HCDs propose abstract comprehensive clusters of empirically related cultural manifestations. In this direction HCDs introduce five major categories: PDI-CD depicting the cultural sensitivity to power statuses; IDV-CD depicting the cultural sensitivity to self-interest and autonomy vs. common interest and social embeddedness; MAS-CD depicting the preference towards performance, assertiveness and mastery vs. consensus, modesty and harmony; UAI-CD depicting the cultural sensitivity to uncertainty and consequent reactions (e.g. standardization vs. ad-hoc resolutions); LTO-CD depicting the importance of long vs. short-term rewards, promoting strong commitment vs. swift adaptation.

These dimensions study numerous related manifestations, both at the individual and collective levels. For instance, in a high PDI-CD culture, individuals tend to give more importance to leaders when making decisions; at the collective level, decisions tend to be more centralized around leaders, leading to higher performance in simpler environments but lower performance in more complex environments (i.e. when the leader is overloaded or lacks expertise).

In addition to these manifestations, two key generic aspects of the influence of culture are particularly impactful in the context of coordination. Culture is mostly expressed in informal decisions. In order to achieve high performance, culture needs to match the coordination mechanisms in place (e.g. rule-based coordination tends to work better in high UAI-CD cultures). Furthermore, individuals tend to fail to coordinate more often when they do not share the same culture.

First Answer for Research Question 2: *How to Model the Influence of Culture on Coordination?*

Two core mechanisms are identified for modeling culture in situation of coordination: values and practices. These two mechanisms are acknowledged to be the core cognitive mechanisms influenced by culture, supporting in turn the occurrence of most of known cultural manifestations.

Several computational models of these mechanisms are available. Nevertheless, these models remain to be applied for replicating concrete manifestations of culture. Furthermore, models of values also need to be further considered in the context of coordination.

In addition to these models, multiple computational models of the influence of culture are available, sometimes modeling situations of coordination. Nevertheless, these models are not appropriate in this thesis. Indeed, these models rely on integrating behavior-oriented cultural manifestations depicted by HCDs. While being adequate for simulation, where the range of decisions is relatively small and known in advance, these models are not applicable for handling the wide range of complex decisions we aim at handling in this thesis.

First Answer for Research Question 3: *How Can we Use Models of Human Culture as a Practical Tool to Support Flexible Coordination in Artificial Societies?*

This question is not directly tackled by former research and, this thesis aims at further exploring this area.

Nevertheless, several clear indications are provided by techniques for managing culture in human societies, which constitute one of our sources of inspiration. They indicate the importance of taking culture in consideration (i.e. making culture explicit in MASs) and making sure that coordination is in line with culture (i.e. taking both aspects into consideration when designing MASs).

As a more philosophical remark, techniques for supporting flexible coordination aim at letting the system and particularly agents manage coordination on the fly. Several general abstract techniques are available (e.g. cooperation), but these techniques discard an aspect that is central when coordinating: defining “harmony”.

From a fundamental perspective, coordinating consists in helping agents to “work together” by “harmoniously” solving their dependences. Thus, when aiming at letting agents managing coordination, we should specify what we mean by “harmony” and thus by “working together”.

This definition of “harmony” or “working together” is precisely what we want to achieve in this thesis by using culture. As a parallel with human societies, culture dramatically influences what individuals consider as “working together” (e.g. what is important, orders from the boss? Rules? Good performance? All being satisfied?) and thus how they behave when trying to work together. If their understanding and expectations do not match, coordination is likely be limited or fail, particularly when trying to handle new situations.



3

Towards a MAS-Compliant Conceptualization of Visible Influences of Culture on Coordination

Chapter Summary

This chapter aims at conceptualizing the visible influences of culture on coordination such that this conceptualization can be related to MASs. In particular, we consider the core conceptual and empirical “clusters” of the influence of culture as depicted by Hofstede’s Cultural Dimensions (HCDs). Based on this theory, we introduce a conceptualization of the influence of culture on coordination that is compliant with the MAS perspective.

As a first step towards answering the research questions set by this thesis, the current chapter aims at conceptualizing the influence of culture on coordination in human societies. In particular, this conceptualization should relate easily to the MAS framework. More precisely, we aim at providing the *core concepts* for explaining the *main patterns* of the influence of culture on coordination (e.g. culture influences leaders towards being directive vs. consultative).

Why not considering more specific conceptualizations? Indeed, such a conceptualization can help reproducing better the influence of culture on coordination (e.g. management by objectives works better in the United States than in France, as stated in Hofstede et al. (2010a)). Nevertheless, aiming for further specification has limited interest. First, because being capable of conceptualizing, replicating and handling the key aspects of the influence of culture on flexible coordination is sufficient with regard to our research goal (showing that culture can be used as a tool for supporting coordination). Second, because we aim at keeping our model simple and practical for sake of coordinating agents. Thus, since modeling in details more extensive conceptualizations is likely to generate more complex models, the gain from expanding conceptualizations is limited.

In order to achieve the goal of the current chapter, we look for key concepts that can be used for explaining the main patterns of the influence of culture on coordination. Certain former research explores this direction, such as Hampden-Turner and Trompenaars

(1993); Hofstede et al. (2010a). In this thesis, we focus on Hofstede's Cultural Dimensions (HCDs) theory from Hofstede et al. (2010a), which has received the most attention with regard to available theories. Each HCD introduces a "variable" (a score) which measures the influence of culture on certain aspects of individuals and societies (e.g. PDI-CD: sensitivity to status; UAI-CD: sensitivity to uncertainty). Each HCD can be conceptually related to numerous concrete cultural manifestations (e.g. PDI-CD is related to the tendency to blindly accept orders and to avoid taking initiatives). These manifestations are correlated to each other (e.g. in a high PDI-CD culture, subordinates are likely to both blindly accepting orders and not take initiatives). In sum, the score of each HCD denotes a range of correlated cultural manifestations to be expected. Therefore, HCDs provide an adequate basis with regard to the goal set by this chapter. We build our conceptualization based on this theory and towards raising MAS-compliant concepts.

The question of the validity of our conceptualization is to be raised. As for any observation-based theory, our conceptualization cannot be "proven exact". Nevertheless, such theories can be given a degree of confidence that indicates the range of application where this theory is relatively adequate. In our case, our conceptualization is tightly coupled with the theory of Hofstede et al. (2010a): this later theory describes the influence of culture in general and our conceptualization focuses on the case of coordination. While not being empirically validated, our conceptualization is streamlined with the theory of Hofstede et al. (2010a) that is empirically validated. To that extent, we can be relatively confident that our theory is conceptualization accurate in comparison with the theory from Hofstede et al. (2010a). This relative accuracy is sufficient with regard to our end-goal: we want to reproduce the core features of the influence of culture on coordination.

In addition of the conceptualization introduced in the current chapter, we relate the type of culture depicted by HCDs with values to be expected in this culture. We perform this step here because we rely on values for modeling the influence of culture on coordination in later chapters. Thus, we need to relate values with visible manifestations as depicted by HCDs in order to validate and use our value-based models of culture for supporting coordination.

Regarding the rest of this thesis, this conceptualization expands the key findings from social sciences about the influence of culture in general as presented in Chapter 2. Then, this conceptualization will be used in following chapters as a basis for building agent-based models of the influence of culture on coordination (Chapter 4); for creating expectations about culturally-sensitive behaviors to be expected within MASs integrating such culturally-sensitive agents, for sake of validation (Chapter 5); for expanding available social science theories in order to explain the individual-to-collective link (Chapter 6); and for sake of determining how culture influences coordination and thus can be used for supporting coordination (Chapter 7).

This section is structured as follows. In Section 3.1, we provide the general MAS-compliant concepts for depicting the influence of culture on coordination as described by HCDs. In Section 3.2, we conceptualize the relation between values and HCDs.

3.1 Conceptualizing the Influence of Culture on Coordination Through Hofstede's Cultural Dimensions

This section aims at expanding HCDs for detailing the influence of culture in the context of coordination by using MAS-compliant concepts. HCDs describe the core variables for characterizing visible manifestations of the influence of culture in general and in more specific context (e.g. at school, at the workplace). We use these descriptions for conceptualizing this influence in the context of coordination. More precisely, given these descriptions, we focus on the context of organizations. Furthermore, we rely on concepts which are or can be easily translated in MAS-related terms. This research is based on one of our articles: Vanh e et al. (2013a).

We do not claim this conceptualization to be exact exact or complete. Nevertheless, this conceptualization introduces the core variables of the influence of culture on coordination. This conceptualization is tightly coupled with descriptions provided by HCDs, which are strongly empirically validated. To that extent, we can be confident that this conceptualization is relatively accurate.

In the following sections, we conceptualize the cultural tendencies expected to be found in an organizational context as depicted by the various values of each HCD. For each extreme score of each HCD, we indicate the type of cultural behaviors to be expected (e.g. in high PDI-CD culture, subordinates tend to comply more easily with instructions from leaders than in low PDI-CD culture).

Each following section is illustrated by a conceptual graph. Nodes of these graphs represent culturally-sensitive manifestations that tend to occur (or not) depending on the type of score this culture would achieve in a given cultural dimension. Some of these phenomena can be related to (possibly causing) the other (e.g. "obedience" is connected to "importance given to leader's opinions"). In that case, these phenomena are connected to a plain link. Some of these phenomena can be negatively related to another (e.g. "making sure that the leader has plain information" and "taking initiatives"), and in that case they are connected to a dotted line. Certain phenomena can be in direct opposition (e.g. power-dependent vs. power-independent status allocation). In that case we represent these two aspects in a single box in order to avoid multiplying links. This representation aims at introducing the core relationships without looking for completeness. In particular, we avoided redundancy and indirect links for sake of readability (e.g. PDI-CD influences both "leaders should care for the group" and "obedience", while the two are also conceptually related to each other).

3.1.1 Power Distance

PDI-CD is strongly related to the cultural importance given to power. In this direction, coordination and organizations in particular, introduce power relationships through delegation structures amongst others (e.g. good performance, expertise can also confer power status in a lesser extent). For the sake of conciseness, this section refers to "leaders" as individuals that are given a formal power status and to "subordinates" as individuals with less formal power, assuming that both are tied by a power relationship.

In high PDI-CD organizations, leaders tend to be considered as superior by subordi-

nates (e.g. being more important, more informed, wiser). Leaders tend to be in charge of managing information and decisions about the group. Subordinates report information to leaders and require from them to make decisions concerning the group (e.g. approval when taking initiatives or for interacting with other subordinates). Furthermore, subordinates tend to be relatively compliant to decisions made by leaders. In high PDI-CD organizations, leaders and subordinates tend to be tied by a strong dependence relationship. Leaders are in charge for “managing and protecting” subordinates.

In low PDI-CD organizations, leaders tend to be given an equal status by subordinates. The influence of formal power on collective decisions is blurred, even if informal influences can still be at play (e.g. giving more importance to experts or good performers). Leaders tend to consult subordinates when making decisions and try to make decisions that are in accordance with the collective agreement. All individuals tend to accept to share the responsibility for collective decisions that are made. Subordinates tend to initiate interactions more easily with each other and take initiatives without requiring approval from leaders. In conclusion, leaders and subordinates tend to be relatively independent from each other.

Influence on Status: In high PDI-CD organizations, leaders tend to formally and informally receive and expect more status from other individuals than in low PDI-CD organizations. In other words, in high PDI-CD organizations, power tends to grant more status (accepted influence on decisions, being more easily supported by individuals with lower subordinates) than in low PDI-CD organizations.

Influence of PDI-CD on Organizational Decisions in General: The degree of PDI-CD influences the likelihood that subordinates accept opinions and influences of leaders. PDI-CD influences the relative weight that subordinates give to their own opinion. Furthermore, PDI-CD also influences how many concessions subordinates are willing to make towards leaders.

As a consequence, individuals in a high PDI-CD culture are more likely to expect that leaders make collective decisions and to blindly comply with these decisions. Conversely, individuals with low PDI-CD cultures are more likely to prefer being heard and more democratic decision processes.

Influence on Interactions Related to Delegation: In high PDI-CD organizations, subordinates are more likely to accept blindly decisions of leaders. Subordinates tend to assume that leaders make adequate, better informed decisions and that they should not be questioned. Thus, instructions from leaders tend to be seen as “orders” that should not be discussed. Conversely, in low PDI-CD organizations, these instructions tend to be seen as “propositions”, letting subordinates deciding on their own whether they can or want to follow them.

As a consequence of these possible interpretations, in high PDI-CD organizations leaders are expected to and thus responsible for carefully checking the correctness and concrete feasibility of their instructions. Otherwise, since subordinates are expected to accept instructions, possibly coordination-damaging situations can happen (e.g. avoidable overload, lower quality, failure). Conversely, in low PDI-CD organizations, leaders can make proposals more freely because subordinates tend to be expected to check whether instructions can be done and to react if they cannot or make counter-proposals. This point is at the core of our simulation presented in Chapter 6.

Influence on Failure Management: In high PDI-CD organizations, subordinates tend to

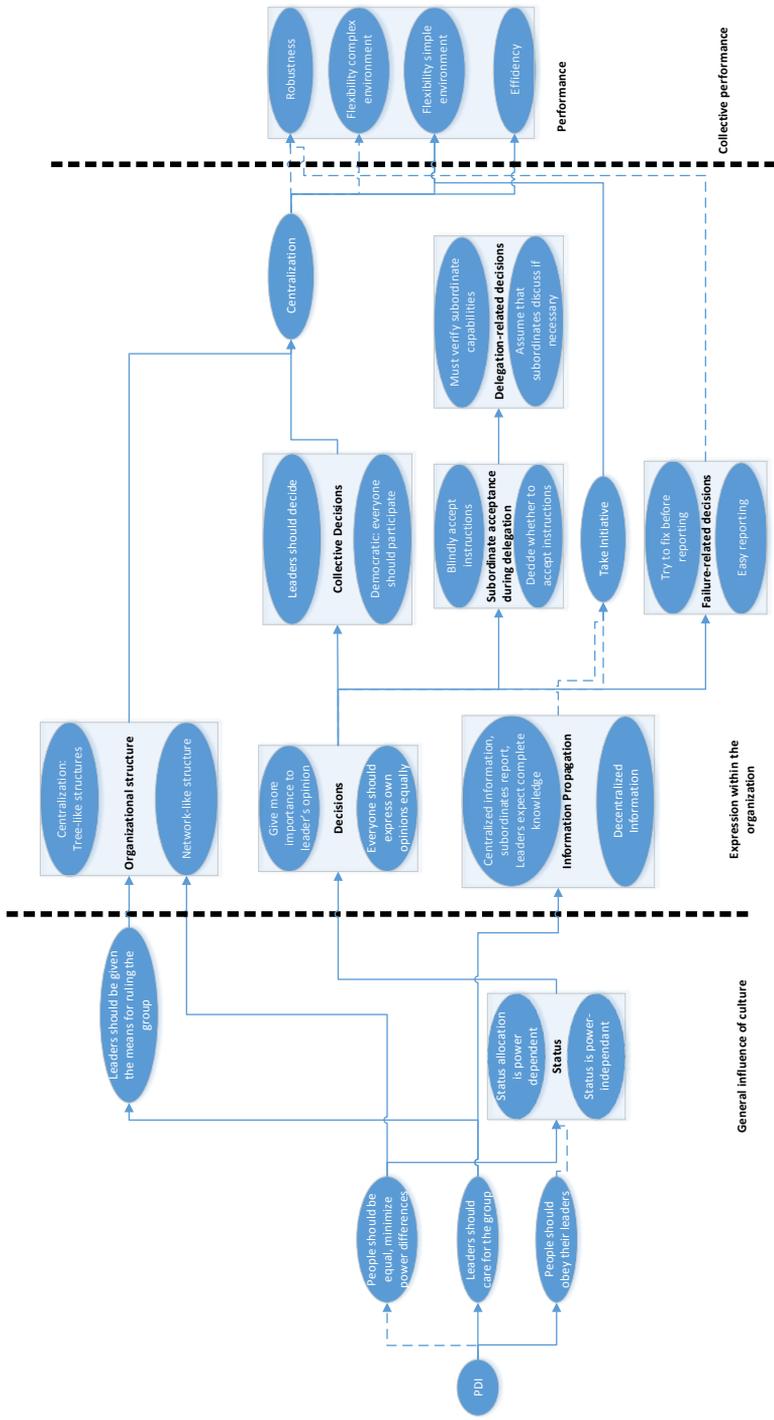


Figure 3.1 – Conceptual graph relating the influence of PDI-CD on organizational behavior

be committed longer to instructions given by leaders and try to repair it locally if possible. Thus, they tend to report failures much later. This decision may also prevent to force leaders to re-plan for fixable failures. This decision is particularly adequate because leaders are in charge of making more decisions for subordinates and thus tend to be overloaded.

Conversely, in low PDI-CD organizations, formal power has less influence on the delay for reporting failures. Thus, subordinates are more likely to report failures earlier to leaders.

Influence on Information Management: In high PDI-CD organizations, subordinates tend to provide extensive information to leaders such that they can make informed decisions. Thus, subordinates are more likely to centralize information about the environment around leaders.

This tendency is less visible in low PDI-CD organizations. Subordinates are more likely to keep private any information that is not useful for the leader (e.g. directly asking for help from a colleague without the leader). Therefore, information tends to be more decentralized.

Influence on Signal Origins (Taking Initiatives): In high PDI-CD organizations, subordinates are less likely to take initiatives. Leaders are expected to plan for the group and subordinates are expected to not conflict with those plans. They tend to require at least the approval from leader, which may prevent to take initiatives due to inherent costs.

Conversely, in low PDI-CD organizations, subordinates tend to be more independent and thus to take more initiatives. Furthermore, they can inform leaders more easily, without impacting on their plans. They can also enter in direct contact more easily with other subordinates, without disrupting the command chain. This can be used for requiring help when taking initiatives.

Influence on Structure: High PDI-CD organizations tend to be structured around tree-like organizational structures. These structures centralize information, decision, delegation, monitoring and failure handling (e.g. machine bureaucracy, simple structure). Such a structure provides concrete organizational aspects driven by a high PDI-CD culture: leaders are provided with information and decision power. In return, subordinates are coordinated with only one leader, preventing conflicts when blindly accepting instructions from multiple leaders.

Conversely, in low PDI-CD organizations, power status tends to be more spread, leading to more decentralized organizational structures (e.g. monitoring and delegating being performed by two different roles). This structure supports more even distributions of power within the organization, preventing a few individuals to be overloaded and to make all decisions. Therefore, structural merging is avoided, leading to more balanced and robust networks (e.g. adhocracies and professional bureaucracies).

Influence on Structure Evolution: Since leaders tend to be given more status in high PDI-CD organizations, they tend to impact more the evolution of the organizational structure. Conversely, in lower PDI-CD organizations, this process is expected to be more democratic.

Influence on Performance: In relating PDI-CD to performance, high PDI-CD organizations tend to achieve higher performance in simple environments and lower performance in more complex environments. Furthermore, PDI-CD tends to support lower robustness. High PDI-CD organizations tend to be more centralized, which is adequate in simple environments (leaders can timely centralize and optimize allocations), but less for more com-

plex environment (receiving and processing information takes too much time and expertise). Conversely, decentralization promoted by low PDI-CD tend, to lead to a reverse tendency: better capability for handling complex environments (complex decisions can be done relatively locally) and higher robustness, but lower results in simple environments due to more limited possibilities for tightly optimizing coordination.

3.1.2 Individualism

At the core, IDV-CD influences what individuals consider as their identity.

In high IDV-CD organizations, individuals tend to connect their identity to themselves as individuals. They see themselves and others as autonomous and independent beings. Everyone should determine what he or she wants and pursues it by his or herself and do not expect to be helped by relatives unless they have an interest in doing it. This desire for autonomy does not prevent from acting in groups, but group cohesion is mostly set by matching personal interests. In high IDV-CD organizations, individuals tend to be independent from each other.

In low IDV-CD organizations, individual identity tends to be more merged with the identity of social groups that they belong to. Individuals tend to consider themselves as deeply embedded within their social contexts (e.g. work group, family, neighborhood). They tend to comply with the “social identity” of their groups (e.g. goals, rules, what is “good” for the group). They tend to create strong informal links, to support and to create mutual dependencies with each other. A key aspect of low IDV-CD is the tendency to adopt mutual protection relationships between individuals and their groups: individuals support their groups and other individuals within their groups and expect in return to be supported and protected by the group.

Coordination in general and organizations in particular, promote interactions between individuals and thus the creation of informal links (e.g. individuals interacting along the workflow, in the same unit). These informal links can enrich formal organizational links. In this section, the categories of individuals with which an individual can interact with are referred to as *groups*. Members of these groups are referred to as *fellows*.

Note that the influence of IDV-CD partly overlaps with PDI-CD, particularly concerning the creation of dependence between individuals or not. Both cultural dimensions are conceptually and empirically related.

Influence on Status: Individuals belonging to similar groups tend to give certain status to each other. IDV-CD influences how much status tends to be given for belonging to the same group. In high IDV-CD organizations, individuals that belong to the same group are not given more status for this reason. Conversely, in low IDV-CD organizations, individuals tend to give and expect more status from fellows.

Influence on Decision: IDV-CD indicates the relative preference towards favoring self-interest over interests of group and fellows. In high IDV-CD organizations, individuals tend to favor self-interest. They avoid making preferences based on their groups or informal relationships.

In low IDV-CD organizations, individuals tend to support the group and fellows (nepotism). This support can take over personal interest. Furthermore, they tend to comply with group norms, which enforces group interest. Individuals tend to prefer acting in groups, with other individuals.

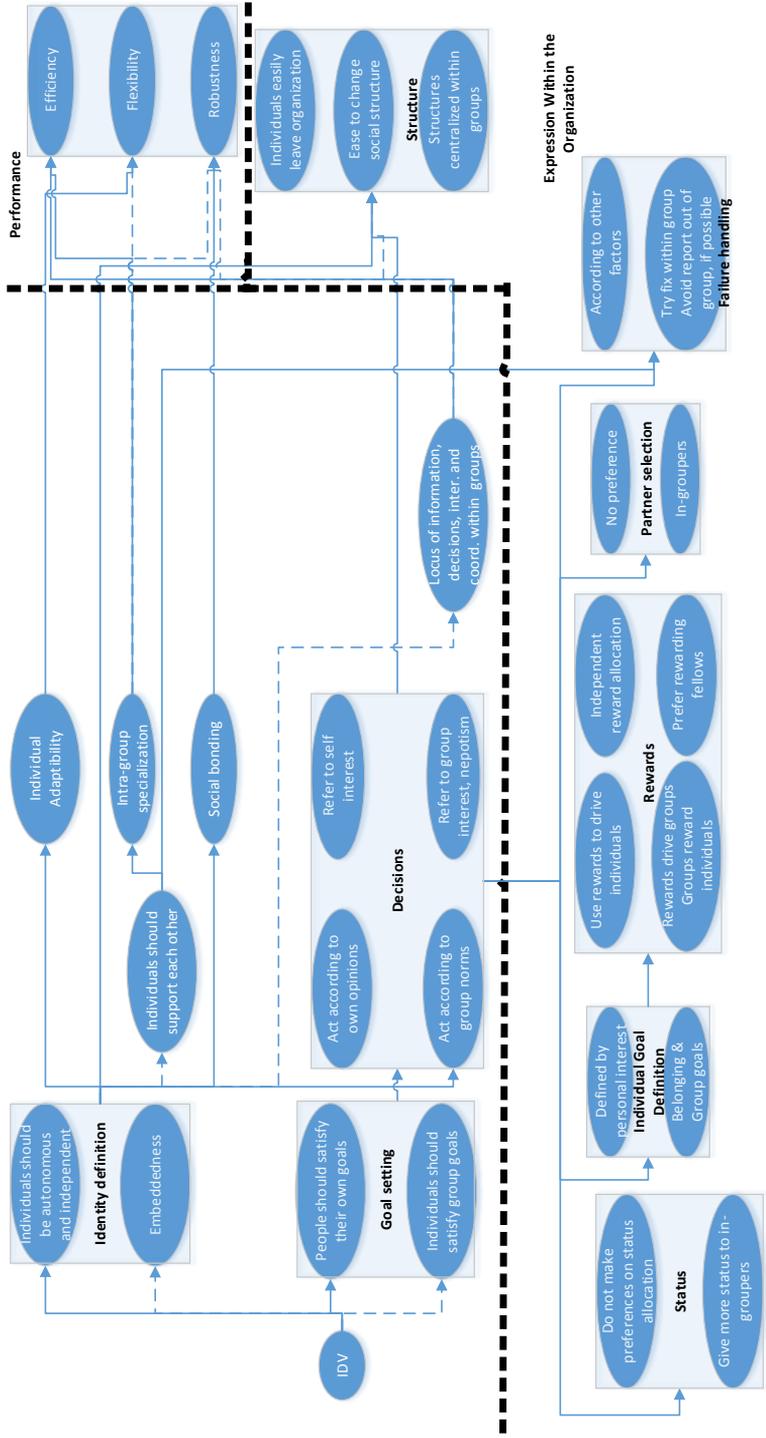


Figure 3.2 – Conceptual graph relating the influence of IDV-CD on organizational behavior

Influence on Individual Goal Definition: Depending on the degree of IDV-CD, individuals may be more or less sensitive to the social context they belong to. In high IDV-CD organizations, individuals tend to pursue their own desires. In such a culture, individuals can be driven by rewards satisfying their own desires.

In low IDV-CD organizations, individuals tend to give more emphasis to goals matching the interest of the group and of other individuals within the group. They are also more sensitive to their social context they want to preserve. Furthermore, in low IDV-CD organizations, group rewards tend to be given more importance than individual rewards. Finally, individuals are more likely to prefer allocating rewards to in-groupers.

Influence on Failure Handling: High IDV-CD tends to have limited cultural influence on the way failures are handled. In low IDV-CD organizations, individuals tend to prefer handling failures within group or to report the responsibility to another group. Failures tend to be “hidden” within group for protecting the image of the group.

Influence on Interactions: High IDV-CD tends to support a preference towards being independent from each other when interacting. Individuals prefer to be capable of solving their tasks on their own. This preference towards autonomy allows individuals to change roles more easily.

Low IDV-CD tends to favor tight interactions within group. These intra-group interactions tend to support local specializations.

Influence on Information Propagation: In high IDV-CD organizations, information is seen as universal. Information can be shared relatively equally with everyone. Group membership does not influence transmission of information. In low IDV-CD organization, locus of information tends to arise within social groups. Individuals tend to share extensively their information within group but less out of group. This exclusivity tends to create a rich communication context that cannot be easily grasped by out-groupers (e.g. “private jokes”, group-specific terms). Furthermore, freely delivering information to out-groupers tends to be avoided since it may contradict the interest of the group.

Influence on Structure: In high IDV-CD organizations, individuals tend to have few restrictions on the range of organizational shapes that can be adopted. In low IDV-CD organizations, individuals tend to prefer to be related to other individuals with similar social contexts (e.g. similar education background). Furthermore, individuals tend to dislike strongly structures allowing the “intrusion from out-groupers” (e.g. being monitored by an external entity).

Influence on Structure Evolution: The structure of high IDV-CD organizations tends to evolve easily. Individuals can be easily replaced and reallocated, since they tend to avoid creating strong dependence links (both affective and operational) with each other. From a general perspective, organizations can adapt more easily to evolving environments.

In low IDV-CD organizations, individuals tend to prefer avoiding evolution of the organizational structure. The main reason is that they create strong bonds between each other, thus they have to re-create bounds in the new group. New-coming individuals are seen as out-groupers, which are given lesser status until being integrated. Furthermore, individuals tend to create specific interaction patterns, supporting local specialization. Thus changing the structure implies breaking a local organization that has to be rebuilt, thus possibly leading to a high performance loss.

IDV-CD influences how organizations are to be used. When scoring high on IDV-CD, organizations change their internal structure to the problems they want to handle. When

scoring low on IDV-CD, organizations tend to look for the problems that can be well handled by their internal structure.

Influence on Performance: High IDV-CD supports relatively high flexibility, due to high individual autonomy, relative independence and ease to adapt the organizational structure on the fly. Conversely, low IDV-CD supports relatively higher efficiency, but lower flexibility and higher robustness. This performance profile is related to the specialization, interdependences and mutual support promoted by a IDV-CD culture.

3.1.3 Masculinity

The core of the MAS-CD can be strongly related to the individual and collective influence of culture on the answer to three dilemmas: “mastery versus harmony”, “equity versus equality” and “assertiveness versus modesty”. High MAS-CD organizations tend to support the first options while low MAS-CD organizations tend to support the last one. As a note, masculinity is not in direct conceptual opposition with femininity. Nevertheless, empirically, both directions are in opposition in practical situations: favoring one is generally done at the expense of the other.

Organizations offer a rich background for expressing these dilemmas. Organizations provide performance indicators for measuring levels of mastery (e.g. number of items produced per minute). Organizations provide rewards to be shared without formalizing how these rewards should be shared (e.g. on merit or equally). Organizations support rich interactions where harmony and consensus can be searched and assertiveness and modesty can be expressed.

Influence on Status: MAS-CD influences the relative status provided by the capability for being an achiever versus the capability for managing positive interactions with others (e.g. everyone is satisfied, mutual respect).

In high MAS-CD organizations, individuals tend to give more importance to mastery when allocating status. Mastery is visible through good performance, rewards and achievements. In low MAS-CD organizations, more status is given to individuals whom are easy to live with, when achieving goals without conflicting with others or when managing to prevent conflicts.

Influence on Decision: In high MAS-CD organizations, individuals tend to prefer to have and show high mastery. This mastery can be shown through visible high performance, by using performance indicators (e.g. quality, quantity). As a consequence, individuals are more likely to be driven by these indicators when making decisions. Furthermore, since high mastery is given a lot of importance, individuals are more likely to engage in high-risk high-reward decisions for showing mastery. High MAS-CD organizations tend to drive individuals towards *competition*, which allows comparing the mastery level between multiple individuals. Finally, individuals tend to prefer equity as a way to measure fairness: higher rewards should be given to people with higher mastery.

In low MAS-CD organizations, individuals tend to prefer avoiding conflicts. In order to do so, they are more likely to make decisions that avoid tensions and conflicts with others. These decisions have numerous manifestations: asking and taking into consideration opinions of others, caring about the well-being of others. Individuals tend to be modest in order not to avoid contradicting others or make false promises that can lead to conflicts. Individuals tend to prefer egalitarian repartition of resources, which avoids raising

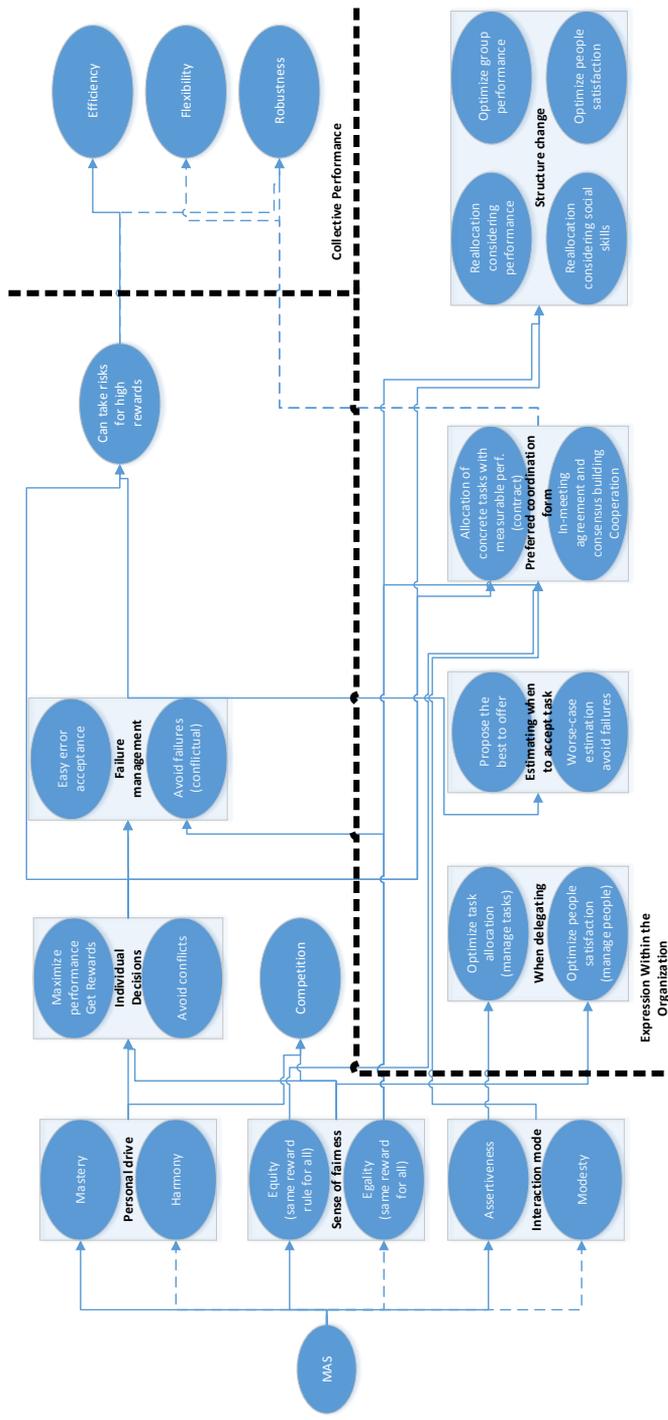


Figure 3.3 – Conceptual graph relating the influence of MAS-CD on organizational behavior

tensions.

Influence on Interactions: In high MAS-CD organizations, individuals are more likely to be assertive. Individuals tend to propose the best they can achieve, even if this proposal may be hard to meet (assertiveness and high-risk high-reward strategy). In general, individuals are more likely to interact through tasks or “challenges” (e.g. contracts as proposed by D’Iribarne (1989)). Individuals are also more likely to reward others for their good performance.

In low MAS-CD organizations, individuals tend to look for consensus (D’Iribarne (1989)). Individuals expect their interest to be taken into consideration in collective decisions. Furthermore, individuals are more likely to cooperate and to support individuals facing issues.

Influence on Delegation: In high MAS-CD organizations, leaders tend to focus on maximizing performance, managing tasks more than managing people. More concretely, leaders tend to optimize task allocation, taking into consideration the mastery of subordinates. This can sometimes lead to uneven distributions (e.g. first-class tasks for good performers). Subordinates are more likely to accept tasks early, being optimistic about their resolution (high risk, high reward).

In low MAS-CD organizations, delegation tends to be achieved through mutual agreement. Leaders care about the satisfaction of subordinates. This can be achieved, for instance, by distributing tasks fairly in order to avoid disagreements and taking into consideration well-being of their subordinates. Subordinates are more likely to committing too early to too many tasks, in order to avoid raising failures, which may lead to conflicts.

This form of delegation is particularly modeled in Chapter 6.

Influence on Failure: High MAS-CD organizations tend to promote high-risk high-reward decisions. Failure is expected for challenging tasks and relatively well accepted. Conversely, failures for simpler tasks show a lack of mastery and tend to be more severely blamed. In both cases, competition drives individuals towards being relatively unsupportive when others face failures.

In low MAS-CD organizations, individuals tend to be more ready to provide support when failures occur. Nevertheless, failing individuals, while not being openly blamed, tend to be pushed away.

Influence on Structure: High MAS-CD organizations tend to deploy performance indicators. These performance indicators are crucial for measuring and comparing the mastery of individuals within the organization. Individuals tend to be driven by these indicators: they should support coordination, otherwise, individuals may adopt coordination-damaging behavioral patterns ¹. Furthermore, high MAS-CD organizations tend to supports equity as a fairness value. Organizations tend to propose reward mechanisms based on individual performance (e.g. tips, grades, titles). For instance, contracts correspond to a structure that is well accepted by high MAS-CD organizations as proposed by D’Iribarne (1989).

Low MAS-CD supports equality as a fairness value, individuals tend to prefer that organizational rewards are shared rather equally between all members of the group. Furthermore, individuals tend to prefer avoiding conflicts, leading to a preference towards flexible

¹The academic system illustrates such a deviation: rewarding researchers based on their number of publications possibly leads to lower publication quality

informal structures enabling to provide a rich ground for finding compromises and preventing conflictual interactions.

Influence on Structure Evolution: In high MAS-CD organizations, individuals displaying good performance are given more status and thus more importance when making decisions related to reorganization. Furthermore, this informal status can lead to conflicts with formal power, which tends to be solved by giving more formal power to good performers (e.g. promote good performers into leaders). This perspective is also supported by equity as a fairness value. Furthermore, when reallocating individuals, empirical evaluations of performance (e.g. high performance for a given task) tend to be given more importance than other traits that are harder to quantify (e.g. informal leadership skills).

Low MAS-CD puts more emphasis on conflict avoidance. Thus, organizational changes are more likely to give more importance to desires of individuals, to social aspects (e.g. matching personalities) or to social skills (e.g. being able to reach consensus).

Influence on Performance: High MAS-CD tends to support higher efficiency, but more limited flexibility and robustness. High MAS-CD individuals and organizations tend to be driven by performance measures, which are generally driven towards efficiency (flexibility and robustness are harder to measure). Furthermore, competition tends to prevent mutual support preventing to adapt to new situations.

Low MAS-CD tends to drive individuals towards consensus and cooperation, which generally leads to overall lower efficiency for finding collective agreement. Nevertheless, gathering everyone's opinion allows successfully coping with a greater range of tasks, increasing flexibility. Finally, failures trigger collective support, which supports higher robustness.

3.1.4 Uncertainty Avoidance

UAI-CD corresponds to cultural sensitivity towards uncertainty. UAI-CD can be conceptualized in considering the causes of anxiety generated by uncertainty and mechanisms that individuals use for lowering or coping with anxiety.

UAI-CD can be conceptually related to three main sources of uncertainty: uncertainty raised by incomplete information about the current situation, uncertainty raised by unpredictability of the future and uncertainty raised because of the unpredictable nature of interactions with other individuals. For each of these sources, specific behaviors can lower uncertainty.

Two main tendencies for handling uncertainty can be visible. In high UAI-CD organizations, individuals tend to create strong expectations that the environment is standard for lowering uncertainty. They act towards keeping low uncertainty through their action (e.g. rely on standardized resolutions, warn others if something unexpected happened). In low UAI-CD organizations, individuals tend to create few expectations about standardization. They act towards getting sufficient information for achieving their goals, supporting ad-hoc adaptive resolutions.

Influence on Status: In high UAI-CD organizations, experts tend to be given more status. Indeed, expertise is an unambiguous sign of capability for understanding, handling a predicting the environment. In addition, individuals who rely on standards also tend to be given status, by lowering uncertainty (for others) about their behavior, about the consequences of their actions and about the state of the environment.

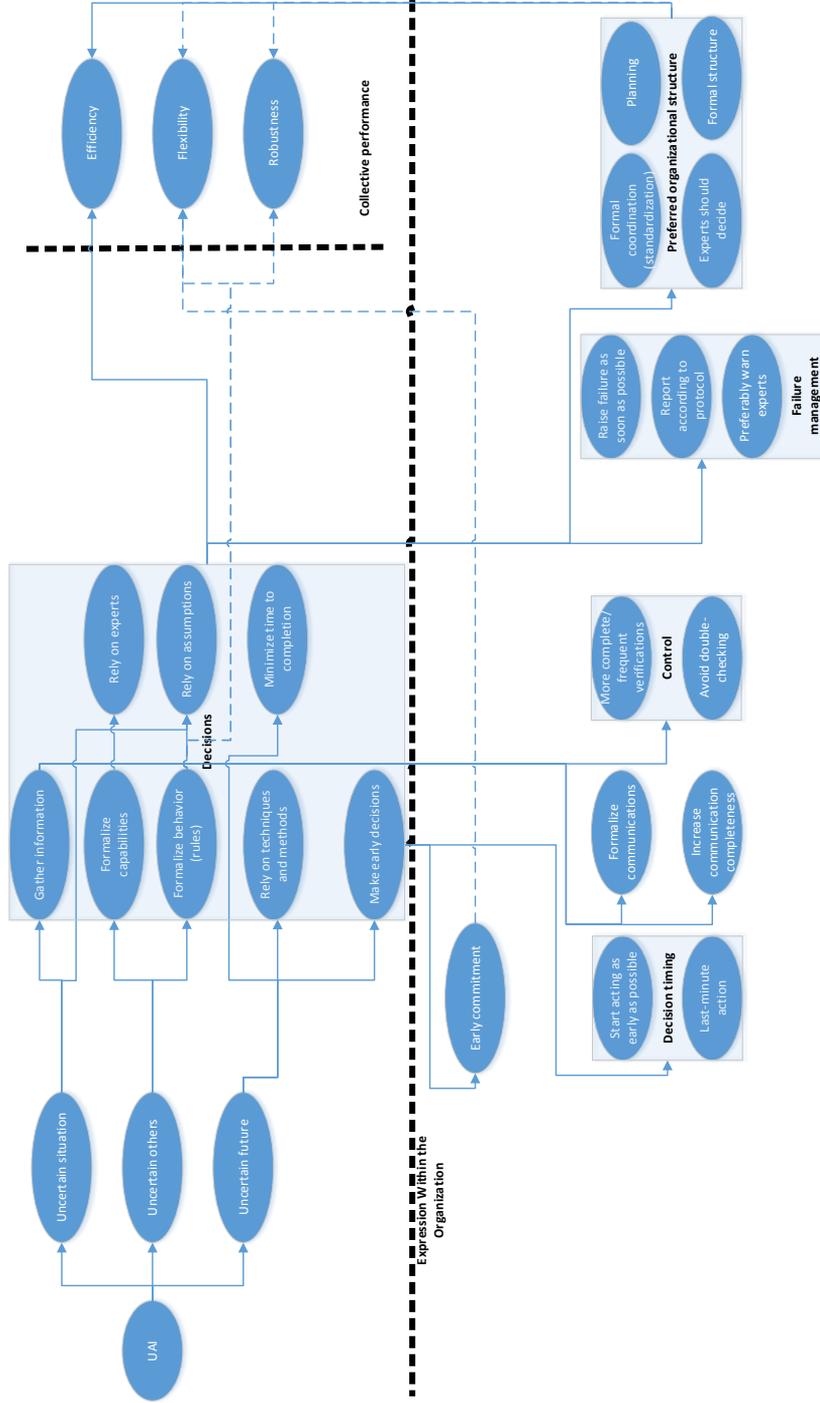


Figure 3.4 – Conceptual graph relating the influence of UAI-CD on organizational behavior

In low UAI-CD organizations, expertise is given less importance. Other aspects such as good performance are then more influential on status.

Influence on Decisions: Uncertainty can arise from three sources: uncertain situations, uncertain outcomes and uncertain interactions. These sources have concrete manifestations within organizations.

Individuals can cope with anxiety caused by *uncertain situations* in two ways. The first solution consists in directly lowering uncertainty by gathering information, through probing, controlling by communicating with others. The second solution consists in lowering anxiety by creating assumptions about the environment (e.g. incoming items are expected to be operational).

Anxiety raised because of *uncertain future* can be linked to the following influences on decisions. First, individuals may be uncertain about the outcome of their decisions. As a solution, they are more likely to prefer relying on planning and on known well-working standards. Second, unexpected events can change current plans, leading to uncertainty (e.g. an extra step has to be performed). This can support early commitment strategies, for achieving plans as early as possible, making sure they are fulfilled. Third, the future can also raise uncertainties because more tasks are expected to arise later. In this case, individuals can adopt strategies for maximizing local efficiency (synchronous cultural dimension from Hampden-Turner and Trompenaars (1993)), possibly at the expense of timeliness.

Finally, anxiety can be raised because of *social uncertainty*. This anxiety is caused by unexpected decisions of other individuals. In order to cope with this uncertainty, individuals can create rules and bring incentives for following them. Individuals can also rely on formal communication in order to limit ambiguities. Furthermore, individuals are more likely to spend more time communicating for avoiding ambiguity (e.g. being informed about plans of others, forming plans together). Individuals are likely to rely on experts who are expected to provide additional certainty.

Influence on Control: High UAI-CD tends to support extensive, frequent and complete control, in particular through formal verifications. Individuals tend to create strong expectations about control, assuming the environment and the work of others to be conform to expectations.

Influence on Interactions: UAI-CD influences the preference towards relying on clear and standardized interactions.

In high UAI-CD organizations, individuals tend to create and rely on explicit rules for solving problems. This standardization has two benefits on lowering uncertainty. First, they provide a guideline for behaving, avoiding anxiety caused by uncertain consequences of actions. Second, they reduce the range of possible outcomes, lowering situation uncertainty for other individuals. When communicating, in high UAI-CD organizations, individuals tend to rely on standardized communication frameworks. In general, individuals are likely to spend effort in clarification.

In low UAI-CD organizations, individuals tend to rely on less rules and to not seek maintaining low uncertainty. Individuals tend to be more adaptive and careful when exchanging tasks and information.

Influence on Delegation: In high UAI-CD organizations, leaders tend to clarify job descriptions as much as possible. They tend to be expected to be capable of answering questions of their subordinates. When coordinating subordinates, leaders are more likely to raise rules and detailed plans.

In low UAI-CD, leaders are less driven towards relying on rules for coordinating subordinates. Subordinates tend to be more expected to figure out by themselves how to solve problems and to adapt to unexpected situations.

Influence on Failure: In high UAI-CD organizations, individuals tend to raise earlier warnings when unexpected situations occur. They are more likely to warn other concerned agents and request support from leaders or experts.

In low UAI-CD organizations, individuals are more likely to try to solve unexpected situations by themselves. They are more likely to require less support from experts and to be more adaptive in general. Nevertheless, remaining failures tend to be recognized later than in high UAI-CD organizations.

Influence on Structure: In high UAI-CD organizations, individuals tend to prefer relying on explicit organizational structures, which helps determining who should be contacted depending on the situation. Leaders are preferably experts in the domain they manage, in order to answer questions of their subordinates. Individuals prefer to be given rules for making decisions. In conclusion, high UAI-CD organizations tend to be based on standardized organizational structures, like machine bureaucracies or professional bureaucracies.

In low UAI-CD organizations, individuals tend to prefer relying on implicit or minimal organizational structures (e.g. simple structure, adhocracies). These implicit structures support the creation of adaptive networks allowing unconstrained interactions.

Influence on Structure Evolution: High UAI-CD organizations tend to create more rules for coordinating individuals. Organizational changes are likely to be reflected in terms of changes of regulations or explicit workflows. Furthermore, experts tend to be given more status, thus more responsibility when changing the organization and more influence on delegation.

Low UAI-CD organizations tend to feature informal and flexible organizational structures. Organizational changes are more likely to happen frequently and to be difficult to handle formally or to control.

Influence on Performance: Considering performance, high UAI-CD organizations tend to achieve higher efficiency, but lower flexibility and robustness. Such a culture supports strong standardization and tight planning, which tends to achieve high efficiency. Nevertheless, unexpected events tend to imply extensive re-planning, re-organization, costs for warning individuals, need to re-consider decision rules.

Conversely, low UAI-CD organizations tend to achieve lower efficiency, but higher flexibility and robustness. Informal organizational structures can be relatively adapted on the fly. Furthermore, individuals are more prepared to face unexpected situations and thus to rely on adaptive resolution strategies. Nevertheless, these adaptive resolutions are less efficient for handling the standard case.

3.1.5 Long Term Orientation

LTO-CD measures the influence of culture on the relative importance given to longer-term consequences of actions on decisions.

In high LTO-CD organizations, individuals tend to prefer to keep a static activity. They tend to focus on performing the same task and slowly improve themselves over time,

through trials and error. Individuals tend to possess stable and reliable social networks, be persistent in what they undertake and accept failures as a learning experience.

In low LTO-CD organizations, individuals tend to be more versatile and change activity and social networks relatively easily. They are more likely to be adaptable, by efficiently learning new activities.

Influence on Status: In high LTO-CD organizations, status tends to be given to individuals capable of making and committing to long-term decisions.

In low LTO-CD organizations, individuals capable of achieving immediate concrete results tend to be given more status. Furthermore, in such organizations, status tends to be given to individuals who reciprocate services, respect rules and avoid failures.

Influence on Decisions: High LTO-CD tends to favor decisions that are expected to lead to better rewards in a long-term future. In general, individuals rely on the expectation of relatively stable physical and social environments. As an example of such decisions, one can find learning, exploring the environment, saving money, sparing resources, creating strong and long-lasting relationships, waiting for advancement within the same organization. Furthermore, high LTO-CD tends to support “pragmatic” reasoning. This form of reasoning links influences of the environment to practical outcomes while leaving out the conceptualization of non-visible aspects. Such form of reasoning tends to evolve slowly in the long-term, when the environment or its understanding evolves. Individuals assume that their understanding is only partial and may be inconsistent with the way others understand the same problem. Thus, individuals may disagree, but this disagreement tends to be accepted.

In low LTO-CD organizations, individuals tend to consider that they may leave the organization any time. Thus, they tend to be driven towards obtaining immediate good performance and rewards. Furthermore, they tend to disregard longer-term consequences outcomes and thus relative decisions, such as creating long-term relationships, exploring possible solutions or spending efforts on learning. Low LTO-CD tends to drive individuals towards relying on decision models that are both relatively easy to set-up, reuse and communicate (e.g. logical thinking).

Influence on Interaction: In high LTO-CD organizations, individuals tend to be committed to their activities (e.g. long-term learning) and their positions. Their interactions are more likely to evolve relatively slowly. Individuals tend to accept to have a partial understanding of their environment and thus to accept disagreements with each other.

In low LTO-CD organizations, individuals are more likely to change quickly their social networks and the type of interactions they conduct depending on immediate needs. Furthermore, disagreement is strongly avoided, for instance in relying on logical thinking for establishing a common truth.

Influence on Failures: In high LTO-CD organizations, failures are seen as useful learning material. Failures are part of a long trial and error learning process. Failures tend to be accepted and can more easily be openly discussed with each other.

In low LTO-CD organizations, failures are seen as logical faults or rule violations. This hardened acceptance is further emphasized by relatively shorter-sighted relationships, which tends to support immediate punishment.

Influence on Learning: High LTO-CD tends to support pragmatic reasoning. This pragmatism tends to support continuous learning that slowly evolves over time.

In low LTO-CD organizations, individuals expect to change activity any-time. There-

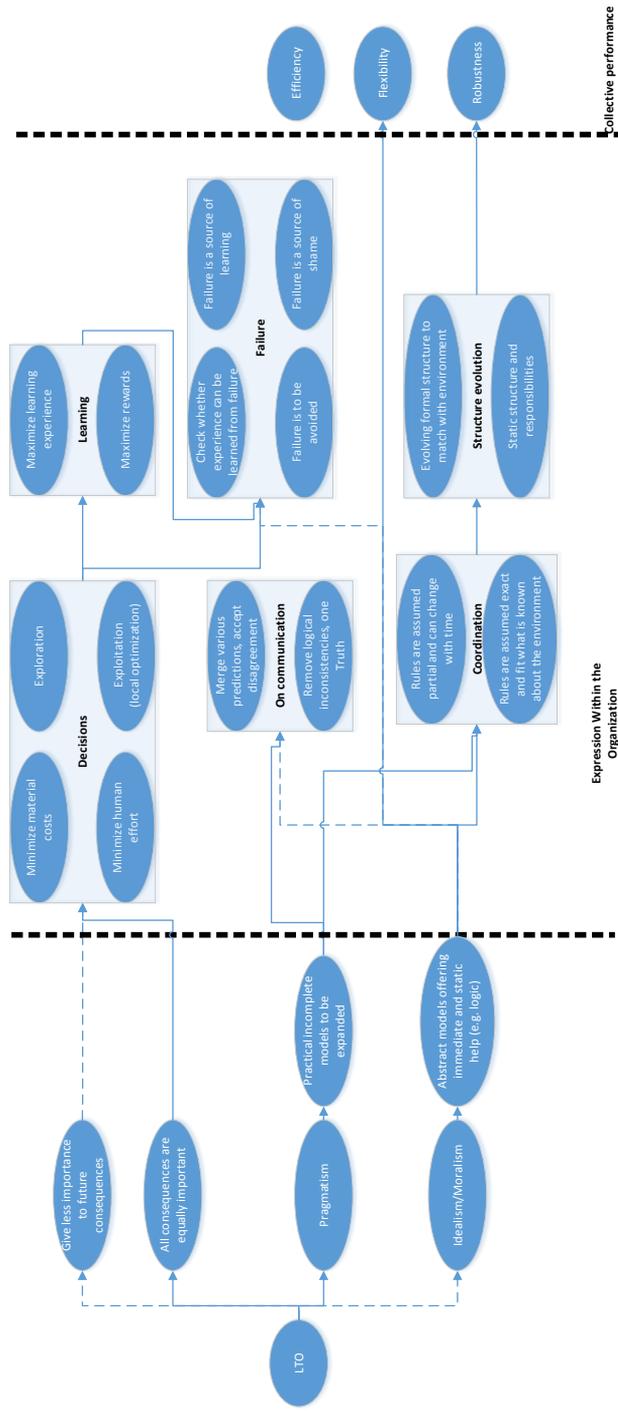


Figure 3.5 – Conceptual graph relating the influence of LTO-CD on organizational behavior

fore, they are more likely to be trained for quickly grasping basic decision rules, but may fail to optimize decisions on the long run.

Influence on Structure: High LTO-CD organizations tend to support learning, making it important to have structures for teaching and gathering information related to learning (e.g. mentorship).

LTO-CD tends to influence rewarding mechanisms. High LTO-CD organizations tend to emphasize longer-term rewards (e.g. savings, stock options of the organization) while low LTO-CD organizations tend to emphasize immediate rewards (e.g. expecting a tip for good work).

Influence on Structure Evolution: LTO-CD tends to influence the desire for organizational adaptability.

High LTO-CD organizations tend to slowly and continuously adapt to their environment. Individuals tend to be more dedicated to a limited range of activities on which they are an expert and to be committed to stable social networks. Nevertheless, pragmatic thinking also supports the evolution of structures for achieving long-term performance. Furthermore, high LTO-CD tends to accept losses of immediate results caused by organizational changes.

Low LTO-CD organizations are seen as more temporary entities, which eventually disappear with time when no longer profitable. Furthermore, changing an organization rarely offers immediate pay-offs without further optimization, which is harder to accept when LTO-CD is low. In addition, in low LTO-CD, learning tends to be more limited which makes decision rules more static, limiting the need for change. Therefore, in low LTO-CD organizations, people tend to have lower incentives towards changing organizations. Instead, individuals are more likely to leave not-so-well-performing organizations and create new organizations or join existing ones.

Influence on Performance: High LTO-CD organizations tend to achieve higher efficiency, lower flexibility and relatively good robustness in stable environments. Long-term learning tends to support high efficiency of when the organization faces repetitive situations. The relative stability of activities and slower learning tend to entail more limited flexibility. Finally, individuals and organizations aim at persisting, supporting high robustness.

Low LTO-CD organizations tend to be more flexible by relying on relatively adaptive logical models, allowing to cope relatively well with unexpected situations.

In conclusion, high LTO-CD organizations tend to achieve higher performance for relatively stable or slowly evolving environments. Low LTO-CD organizations tend to achieve higher performance with highly evolving or unpredictable environments.

3.2 Relating Value-Systems and Hofstede's Cultural Dimensions

In this thesis, values are the core entry-point for integrating the influence of culture on coordination within MASs. As will be justified in Section 4.1.1, values are a core “generative” aspect of culture. Theories of culture such as Hofstede et al. (2010a) put values at the center of culture and thus most of cultural manifestations result from this influence.

Value	HCD
Assertiveness	High MAS-CD
Autonomy	High IDV-CD, low PDI-CD
Conservatism	High PDI-CD, high UAI-CD
Harmony	Low MAS-CD
Mastery	High MAS-CD
Modesty	Low MAS-CD
Obedience	High PDI-CD
Openness to change	Low PDI-CD, low UAI-CD

Table 3.1 – Summary of the conceptual and empirical relationships between the type of values supported by culture and type of Hofstede's Cultural Dimension (HCD) scores

The problem of values is that they are not connected in detail with concrete manifestations. As will be further detailed in Section 4.1.2, values are complex decision mechanisms. Understanding in details their relation with concrete (cultural) manifestations is still an open problem.

This issue is important with regard to this thesis: we need to relate the influence of culture on values with concrete cultural manifestations. This relation is to be used in the following chapters: we need to use it for relating the influence of culture on value-systems to rich conceptual and empirical basis (e.g. studies, cross-cultural comparisons, examples) of concrete manifestations of the influence of culture.

In order to overcome this issue, this section conceptually and empirically relates the values we rely on in this thesis to the type of culture they are promoted, as depicted by HCDs. Thus, we can indirectly relate these values to the extensive conceptual and empirical basis of concrete manifestations of the influence of culture as depicted by HCDs.

In this section, we focus on the following values: assertiveness, autonomy, conformism, harmony, mastery, modesty, obedience and openness-to-change. This list contains all the values that will be modeled in the following chapters of this thesis. Nevertheless, this work can be expanded in relating more culturally-sensitive values (e.g. as proposed for instance by Hofstede et al. (2010a); Rokeach (1973); Schwartz (1999)) with HCDs. These relations are summarized in Table 3.1

Assertiveness: Assertiveness is at the core of the description of high MAS-CD in Hofstede et al. (2010a).

Autonomy: Ng et al. (2007) empirically and conceptually relate HCDs to clusters of values from Schwartz. IDV-CD is related to affective autonomy, intellectual autonomy and egalitarianism. IDV-CD is also negatively correlated to PDI-CD. Affective autonomy and autonomy in general can be expected to be negatively related to low PDI-CD. From a conceptual standpoint, a low PDI-CD culture tends to decouple leaders and subordinates. Thus, low a PDI-CD culture is likely to be related to autonomy.

Conservatism: Ng et al. (2007) empirically relates conservatism to high PDI-CD. From a conceptual perspective, conservatism is related by Schwartz (1999) to conformity (obedience), tradition and security. Obedience is a central part of descriptions of the high PDI-CD in Hofstede et al. (2010a). Tradition and security conceptually relate to high UAI-CD. Thus, conservatism relates to high PDI-CD and high UAI-CD.

Harmony: Harmony is conceptually related to the search for consensus that characterizes low MAS-CD cultures. Furthermore, in Schwartz's Cultural Clusters from Schwartz (1999), harmony opposes mastery, which characterizes high MAS-CD. Thus, harmony is likely related to low MAS-CD.

Mastery: Mastery is empirically related to MAS-CD as indicated in Hofstede et al. (2010a); Ng et al. (2007). This conceptual relationship is also direct with assertiveness, which is itself related to high MAS-CD.

Modesty: Modesty is at the core of the description of low MAS-CD in Hofstede et al. (2010a).

Obedience: Obedience is part of the conservatism cluster of values that is empirically related to high PDI-CD. Furthermore, obedience is explicitly used throughout the description of the PDI-CD in Hofstede et al. (2010a). Thus, obedience is likely related to high PDI-CD.

Openness to Change: Openness-to-change can be related to low PDI-CD and low UAI-CD. Openness-to-change encompasses self-direction and stimulation values, which can be related both with autonomy and creativity. Furthermore as presented in Schwartz (1999), openness-to-change is empirically and conceptually opposed to conservation group of values. This group of values is related to high PDI-CD and high UAI-CD. Thus, openness-to-change is likely a low PDI-CD and low UAI-CD value.

3.3 Chapter Summary

This chapter introduces a MAS-compliant conceptualization of the influence of culture on coordination. This conceptualization is based on HCDs, which introduce five conceptually-coherent variables for capturing the core correlations between concrete influences of culture on coordination. In addition, we relate these tendencies and manifestations with values, which are a core mechanism of the influence of culture and are used in the following chapters.

In conclusion, this chapter provides a general conceptual basis that relates the influence of culture on values and cultural manifestations to expect in specific contexts. This conceptual basis aims at being practical. While not being exact or complete, this conceptual basis is sufficient for creating clear general expectations about what kind of behaviors to expect from a given culture. For instance, if the value of "obedience" is culturally very important, one can expect subordinates to comply with instructions of leaders, leaders to centralize information and relatively limited robustness in cases where the leader cannot be reached.

Regarding following chapters, the conceptualization presented in the current chapter offers a basis that will be used in the rest of this thesis. In Chapter 4, this basis will be used for the sake of modeling the influence of culture. In Chapter 5 and Chapter 6, this basis will be used for checking whether our model of culture can replicate human-like phenomena. In addition in Chapter 6, this basis will provide examples and relationships for supporting our investigation of the individual-to-collective link of the influence of culture in situation of coordination. Finally, the current chapter indicates how culture influences coordination. These influences will be used as a indicators for determining how to "handle" cultures for supporting coordination in Chapter 7.

Answering Research Question 1.a: *How Does Culture Influence Individual Decisions Regarding Coordination in Human Societies?*

This chapter, through HCDs, introduces multiple conceptually-central concrete examples of the influence of culture on individual decisions in situation of coordination.

To re-introduce them briefly, we considered five types of influence. PDI-CD indicates the influence of culture on the importance of power relationships when making decisions. IDV-CD indicates the influence of culture on the relative preference towards being an autonomous self-driven and self-supportive entity versus being embedded within, supporting and being supported by social groups. MAS-CD indicates the influence of culture on relative importance between maximizing performance (e.g. competition) or seeking conflict-less interactions (e.g. cooperation, consensus). UAI-CD indicates influence of culture on the sensitiveness towards avoiding uncertainty and the strategies that are deployed (or not) adopted for handling uncertainty. LTO-CD indicates the cultural tendencies and concrete behaviors towards looking for long-term rewards (e.g. commitment to long-term learning and relationships) or obtaining quicker immediate rewards (e.g. adaptability, easy-to-learn communication frameworks).

These types of influence match features of culture that have been recognized in previous chapters. Individuals expect other individuals to *share* and thus be influenced by the same culture when making decisions. Individuals tend to rely on culture for *creating expectations* about decisions of other individuals. For instance, in high PDI-CD organizations, both leaders and subordinates expect that subordinates should rely on leaders for making decisions. These expectations can in particular support reactions leading to *matching decisions*. For instance, in high PDI-CD organizations, leaders are culturally-driven towards making decisions for subordinates. Finally, the influence of culture is *context-dependent*, in particular, it can be enabled or not by coordination mechanisms. For instance, PDI-CD manifestations of culture are hardly expressed in flat organizations.

As an important notice, HCDs only indicate indirect influence of culture. HCDs investigate visible influence of culture on individual decisions. Nevertheless, this influence results from more direct influence of culture on individual decisions. These more direct mechanisms are introduced in Chapter 4.

Answering Research Question 1.b: *How Does Culture (Indirectly) Influence Collective Outcomes Related to Coordination in Human Societies?*

This chapter relates the key cultural influence depicted by HCDs to collective performance to be expected in situation of coordination.

The type of cultural influence depicted by high score on the PDI-CD (or “high PDI-CD”) tends to lead to high performance in simple environments, low performance in more complex environments and low robustness in comparison with the type of cultural influence depicted by a low score on PDI-CD (or low PDI-CD). High IDV-CD is related to more adaptive coordination, but with less robustness in case of trouble, than low IDV-CD. High MAS-CD is related to higher efficiency, lower flexibility and lower robustness than low MAS-CD. High UAI-CD is related to higher performance in static and predictable environments in comparison with low UAI-CD, but lower performance in evolving or unpre-

dictable environments than low UAI-CD. High LTO-CD is related to better performance in slowly (vs. quickly) evolving environments than low LTO-CD.

As a general remark, even if coordination mechanisms in place influence performance, the influence of culture on collective outcomes appears to be relatively independent from the coordination mechanism.

Answering Research Question 1.c: How Does the Influence of Culture on Individual Decisions Lead to the Influence of Culture on Collective Outcomes in Human Societies?

The key aspects of the influence of culture on coordination depicted by HCDs provide multiple examples of relating the influence of culture on collective outcomes to the influence of culture on individual outcomes. While Hofstede et al. (2010a) do not propose a theory for relating the two, these examples are worth of interest. We build such a theory in Chapter 6. These examples are useful illustrations of this theory.

In high PDI-CD organizations, individuals tend to centralize information around leaders, supporting high performance in simpler environments when they can handle decisions. In low PDI-CD organizations, individuals tend to make decisions locally. This local-ity prevents possible optimizations provided by centralization but also prevents the introduction of a bottleneck that is overloaded when decisions become too complex.

In high IDV-CD organizations, individuals tend to be independent from each other. They can easily change their social context, supporting collective adaptability. Nevertheless, individuals tend to be self-driven, which can prevent mutual support in case of problems, supporting lower robustness. In low IDV-CD organizations, individuals tend to create strong links with each other and to adapt to each other, making more difficult to reorganize agents. Furthermore, they are more likely to support each other in case of problems, supporting higher robustness.

In high MAS-CD organizations, individuals tend to be driven by achieving high performance, according to performance indicators. They are more likely to focus on their tasks without deviating for sake of optimization and compete with each other. This behavior supports high efficiency, but relatively low flexibility and robustness. In low MAS-CD organizations, individuals are more likely to make compromises for making sure that everyone is satisfied. These compromises can be used for enhancing mutual support (high robustness), allowing to achieve a wide range of goals (high flexibility), but with additional costs for finding consensus and for pleasing everybody (moderately low efficiency).

In high UAI-CD organizations, individuals tend to maintain low uncertainty through reliance on standardization. This standardization is adequate for well handling expected situations at the expense of heavy communication for avoiding uncertainty in case of unexpected situations. In low UAI-CD organizations, individuals expect the situation to be uncertain. They locally re-create a sufficient degree of certainty for achieving their goal. This can create ad-hoc resolutions, leading to further uncertain output.

In high LTO-CD organizations, individuals tend to aim for long-term goals. These goals rely on relatively stable social and environmental context, supporting good performance for slowly evolving environments. In low LTO-CD organizations, individuals tend to aim for immediate rewards. They tend to be more adaptive, making them adequate for handling dynamically-evolving environments and social contexts.

As an overview, which is at the core of our theory from Chapter 6, all these examples rely on mutually-matching decisions and thus in coherent and coordination supportive interactions: centralization vs. decentralization, independence vs. inter-dependence, competitive vs. consensual, maintaining low uncertainty or locally managing uncertainty, stability vs. adaptability. These interaction patterns act as “Nash Equilibria”. Culture supports locally-optimal interaction patterns. Individuals who deviate alone from the culturally-supported decisions can expect to lower their own welfare (e.g. mismatching expectations, incoherent resolutions, blame from others), but also collective performance. As highlighted by Hofstede et al. (2010a), culture appears to solve social dilemma. This resolution of these dilemma then drive individuals towards culturally-supported interaction patterns.



4

Replicating the Influence of Culture on Coordination within MASs: from Theories to Models

Chapter Summary

This chapter aims at replicating the core aspects of the influence of coordination within MASs. In order to do so, we first determine the key aspects of culture that (supportively) influence coordination in a wide range of environments (Research Question 2.a). Then, we model these aspects such that they can be integrated within MASs (Research Question 2.b).

In order to achieve the end-goal of this thesis, we need to show whether and how the influence of culture on coordination as observed in human societies can be *operationally reproduced* within artificial societies. More precisely, given the goal of this thesis, we only need to show that the *core aspects* of this influence can be reproduced¹. More precisely, we focus on core aspects of the influence of culture that are expected to influence supportively coordination and that add to existing coordination techniques (e.g. norms, organizations).

Three steps are to be performed for achieving this goal. First step, we need to determine at least a *key mechanism* of culture that influences coordination in human societies (theory) that can add to existing coordination techniques. Second step, we need to show that this key mechanism can be operationally reproduced within MASs (model). Third step, we need to check whether this mechanism replicates the core properties of the influence of culture on coordination (validation).

The first step is handled by exploring theories of culture with the aim of revealing (new) mechanisms of the influence of culture on coordination. Given the goal of this thesis, we

¹In this thesis, we do not aim at replicating all mechanisms of the influence of culture. We only aim at replicating the key aspects of the influence of human cultures on coordination, such that these aspects can be handled in order to coordinate MASs. In particular, we aim at relying on *practical* models. To that extent, we accept to discard some aspects of human culture that can support flexible coordination if their benefit is minor in comparison with the implementation costs for integrating them (e.g. greeting by bowing or shaking hands). Instead, we focus on modeling the core aspects of culture that influence the most the flexibility of coordination.

do not need to explore all possible mechanisms, but only look for the key mechanisms. As a spoiler, the key mechanism of the influence of culture on coordination we found out originates from the influence of culture on individual decisions. More precisely, culture influences certain cognitive mechanisms that influence individual decisions that in turn influence coordination.

The second step is handled in modeling the mechanism that is revealed by the first step. Given the goal of this thesis, we only need to build a model that is capable of reproducing the core features of the influence of culture on coordination. No need for completeness nor exactness, while practicality is an important aspect. Given the mechanism we found, we perform this step in proposing an architecture for building models of the cognitive mechanisms that are sensitive to culture. Then, we propose an agent decision architecture that integrates these culturally-sensitive cognitive mechanism and that is appropriate for handling the situation of coordination.

The third step (validation), is not performed in this chapter. This step requires showing evidence that the mechanisms we modeled accurately replicate their human counterpart. This step will be performed through simulations in Chapter 5 and Chapter 6.

Regarding other chapters, the model of the influence of culture presented in the current chapter relies on the general background describing culture and coordination from in Chapter 2 and our conceptualization of the influence of culture on coordination from Chapter 3. In following chapters, this model will be tested against expectations depicted by social sciences in Chapter 5. This model will be used for building simulations in order to better explain the individual-to-collective link of the influence of culture on coordination in Chapter 6. Finally, this model and underlying core mechanisms of the influence of culture on coordination illustrate that aspects of culture that support flexible coordination can be integrated within MASs. Based on this, we study in Chapter 7 how such a model of culture can be used for coordinating agent societies.

This chapter is structured as follows. First, we investigate what are the key mechanisms of the influence of culture on coordination in human societies in Section 4.1. Then, we propose an architecture for modeling culturally-sensitive mechanisms in Section 4.2. Next, we propose a culturally-sensitive decision architecture in Section 4.3. Finally, we illustrate that this model can be used for replicating a wide range of culturally-sensitive decisions in introducing a range of culturally-sensitive and coordination-influencing decision aspects in Section 4.4.

4.1 Key Mechanisms of the Influence of Culture on Coordination in Human Societies

In order to replicate the influence of culture on coordination, we need first to understand better the core mechanisms of this influence. More precisely, with regard to our goal, we only need to determine a *core* mechanism. This mechanism should be sufficient for replicating the core properties of the influence of culture on coordination (we want to keep our model practical).

Existing theories do not directly point to coordination-specific mechanisms in details. Nevertheless, existing theories as introduced in Section 2.2.3, propose two generic mechanisms for explaining the influence of culture in general: values and practices. The influ-

ence of these mechanisms can then be considered in the context of coordination.

In the current section, we aim at determining whether these general mechanisms influence coordination. More precisely, we are particularly interested in determining which mechanisms can be used for *supportively* influencing coordination in a wide range of environments. These mechanisms are then modeled in Section 4.2 in order to integrate them within MASs.

This section is structured as follows. Section 4.1.1 determines the most important mechanism of culture that can be replicated in MASs for replicating the influence of culture on coordination, in line with the goal set by this thesis. Then, in Section 4.1.2, we introduce more details about how this mechanism works.

4.1.1 Determining the Key Mechanisms of Culture to Focus On

In this thesis, we aim at replicating the core aspects of the influence of culture on coordination. Given the amount of possible situations of coordination, we need to replicate the core mechanisms of the influence of culture on coordination that provide high genericity.

Hofstede et al. (2010a) indicate that the core mechanisms of the influence of culture lie within people's mind. This theory indicates that most of the influence of culture can be boiled down to two cognitive mechanisms: *values* and *practices*. While not being exact, these two mechanisms can be used for explaining most of cultural behaviors and collective consequences of culture. In the current section, we investigate whether these two general mechanisms can be applied with regard to the purpose of this thesis.

The purpose of this thesis raises multiple *criteria* to take into consideration for determining which mechanisms to focus on. More precisely, selected mechanisms should:

1. (Supportively) influence coordination
2. Be influential in a wide range of environments
3. Add to existing coordination techniques

Values (further introduced in Section 4.1.2) fulfill these criteria. Values appear to match well criterion (1): they support coordination. They provide a general context, indicating what is "good" and thus important for individuals. Individuals sharing values (particularly, values supported by culture, such as obedience, autonomy, timeliness) tend to agree better about what are "good" outcomes. Values match criterion (2): they are influential in a wide range of environments. Values can be particularly abstract and generic (e.g. obedience can influence decisions for any leader-subordinate interaction). Thus, they can support many decisions in a wide range of environments. Values also match criterion (3): they add to existing coordination techniques. In considering coordination mechanisms from Section 2.1.4, the closest coordination mechanism is the "incentive-based" coordination (e.g. game-theoretical settings), but this mechanism remains different.

Practices (introduced in Section 2.2.3) fulfill certain of these criteria. As a brief reminder, practices can be seen as cultural behaviors (rituals, e.g. greeting by bowing or shaking hands), important people to follow (heroes, e.g. Napoléon, James Bond) practical and visible symbols (e.g. flags, anthems). Practices match criterion (1). In a strict sense, practices are considered to be signs of recognition, generally without technical interest. Nevertheless, practices feature symbols and rituals, which can be easily translated into

coordination-supporting mechanisms (e.g. ontologies and protocols). Practices partly match criterion (2). Practices tend to be relatively specific for a given context. Furthermore, their possible influences are relatively limited to pre-defined functions. For instance, “greet” can be done in many environments but the range of possible influences is limited. Finally, practices fail to match criterion (3). Symbols and rituals featured by practices are already well handled by existing coordination techniques (e.g. ontologies, protocols). Heroes are a relatively unexplored area for achieving coordination (e.g. learning from imitating an agent that is “important”). Nevertheless, this direction seems to have only a very limited impact on the support provided by culture on coordination.

In conclusion, values are a core mechanism of the influence of culture on decisions that can influence decisions in a wide range of coordination situation and which add to existing coordination techniques. Furthermore, values are shown to be at the core of culture, as depicted in Figure 2.2. They are considered to have the highest impact on individual decisions in comparison with other cultural aspects and to drive the occurrence of other aspects of culture.

As a consequence, the rest of this thesis focuses on the influence of culture on coordination in considering how culture influences values that influence in turn coordination. This section answers Research Question 2.a.

4.1.2 Further Introducing These Mechanisms

Before modeling the influence of culture on values, we need first to conceptualize better the core aspects of values and how culture influences values. Values, particularly when they are supported by culture, are a rich and complex cognitive aspects that influence numerous decision aspects. Theoretical foundations of these decision aspects are further introduced in this section.

Principles of Values

Values are a rich cognitive mechanism extensively studied in Miceli and Castelfranchi (1989).

At the core, values are a special form of beliefs (things assumed to be true by individuals, which can be revised). Values have two uses. First, they indicate what is inherently “good” (e.g. honesty, beauty, respect, cleanliness are “good”). Second, each of these “inherently-good” things provide a perspective for *evaluating* other things (e.g. object, idea, decision). For instance, “stealing” is bad according to the value of “honesty”.

Values have numerous features. They can be more or less *abstract* (e.g. “brushing teeth” can be a value as well as “hygiene”). In this thesis, we focus on abstract values, capable of handling a wide range of environments. Values are *generally but not necessarily correct*. Indeed, values are generally based on features that are “good” or “good for doing something”, often learned from family or trial and errors (e.g. being punished when lying). Thus, they indicate an aspect of something that has led to good results so far. However, values may not necessarily indicate what is the best with regard to a given objective utility (e.g. honesty can incite to report when a delay occurs, even if there is no benefit for doing it). For similar reasons, values are in general relatively accurate but are *not rationally grounded*: additional rational reasoning (which tends to be more expensive) can some-

times show that what is “good” according to a value is actually not desirable. In expressing it with computer-science terms, values can be seen as *heuristics* for evaluating things. Values allow avoiding computationally-heavy rational deliberations in looking for features of the thing to evaluate, which are believed to be generally “good” or “bad”. The “goodness” of these features provides a general idea of the interest of the thing to evaluate.

Values specifically apply to decision-making as depicted by Miceli and Castelfranchi (1989). They can be used as heuristics for making decisions. Individuals can use values for determining whether a decision is “good” or “bad” (e.g. this decision implies lying and thus is “bad” according to “honesty”). The way values influence in detail cognitive processes is very complex and goes beyond the purpose of our models, so are not detailing this aspect here. The core property of the influence of values on decisions to remember for achieving the goal of this thesis is that, when individuals face multiple alternatives for a decision that cannot be rationally distinguished, individuals tend to *prefer value-supported alternatives of decisions*.

Values have several properties that are particularly interesting with regard to the goal of this thesis. In considering both abstractness and the heuristic power, values allow going *beyond rationality*. Values offer irrational yet relatively-accurate and quick-to-evaluate heuristics for making decisions. Furthermore, this aspect-based decision-making, particularly with abstract values, makes that values can support decisions for a very wide range of environments, when rational processes are limited by the lack of information. Furthermore, values still *combine with rationality*. Values take over rational decision processes only when the latter is not sufficient for making decisions. Thus, using values combined with traditional rational deliberation mechanisms offers extra decision support for handling extreme cases.

Value-Systems

Individuals possess multiple values. These values can sometimes contradict each other. For instance, making the decision whether to report a not a fault of my boss. Both alternatives of this decision are hard to distinguish rationally and are supported either by honesty or by obedience. Reporting the failure is evaluated as “good” by honesty and “bad” according to obedience while not reporting the failure is evaluated as “bad” by honesty and “good” according to obedience. Which one to select?

According to authors such as Beck and Cowan (1996); Miceli and Castelfranchi (1989); Schwartz (2006a), individuals *internally order* their values, referred to as the *value-system*. Certain values are *more important* than others. This importance impacts decisions when individuals face value-based dilemmas. In our dilemma, if obedience is more important than honesty, then the individual is more likely to not report the failure. Value-systems do not necessarily solve all possible cases and a gray area can occur (e.g. having to cover a sufficiently important fault), but are still useful in many cases.

In general, individuals possess *coherent* value-systems. They tend to have values that tend to not conflict so much with each other (e.g. obedience and discipline are more often related than obedience and self-direction).

Value-systems have an important impact on behavior. They support *consistency*: values drive decisions and behaviors of individuals towards displaying qualities that this individual considers as important (e.g. this person is honest and disciplined). Since values

evolve relatively slowly, the big features of the behavior of this individual remain the same (e.g. anytime, this individual hardly lies). This consistency is crucial for creating *expectations* about the behavior of this individual.

The Influence of Culture on Values and Value-Systems

Culture is known to influence value-systems as described in Hofstede et al. (2010a); Schwartz (2006a). In this thesis, we refer to values supported by cultures as *cultural values*.

Culture conveys a set of *shared values*. More precisely, people sharing the same culture tend to share similar set of values. *Slight variations* can occur in the way a given value is interpreted by different individuals (e.g. blind compliance can be seen as “good” according to the value of obedience for certain individuals but not for others). In addition, individuals sharing the same culture tend to *share similar value-systems*: individuals give a relatively similar importance to values.

Culture particularly promotes *socially-oriented values*. For instance, Schwartz (2006a) empirically found 10 types of values: stimulation, self-direction, universalism, benevolence, conformity, tradition, security, power, achievement and hedonism presented in Figure 4.1). This social orientation, in combination with the fact that values are shared and impact decisions is crucial in the context of coordination. First, because cultural values support the creation of expectations, indicating what to expect from others. Second, because cultural values support a collective agreement towards certain social patterns. Thus, cultural-values indirectly support the emergence of *institutions* (e.g. norms, organizations) that are in line with values and thus relatively well collectively supported (e.g. laws for protecting workers from power abuse in high PDI-CD cultures). Likewise, values can be used for debating and justifying decisions as presented in van der Weide (2011).

In going a step further in this direction, *values conveyed by culture are not accurate*

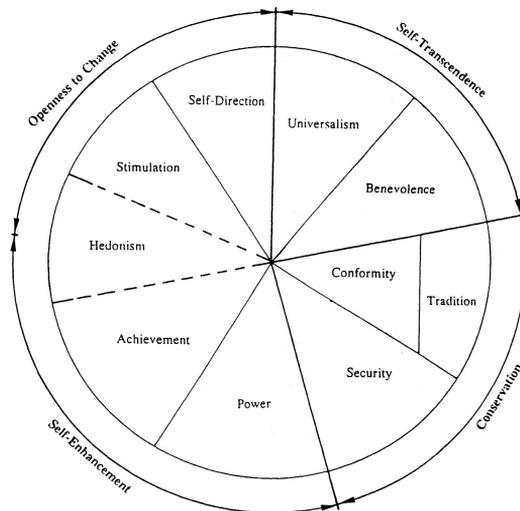


Figure 4.1 – Major types of values according to Schwartz (2006a)

per se. In general, values supported by culture are relatively “accurate” (a value actually indicates what is “good”). They can even become accurate because they are shared. For instance, if obedience is culturally shared then obedience becomes good because in such a culture, disobedient individuals are punished (e.g. being blamed for performing “bad” or unexpected behavior). Nevertheless, “obedience is good” is not always correct, particularly when “obedience” is not a well-supported cultural value (e.g. being told to be a “bootlicker”, to lack initiative).

Given the context of this thesis, we can consider whether values are necessarily adequate for supporting coordination. Research such as Beck and Cowan (1996); Dawkins (1976) suggest that, the answer is no. Cultural values are not necessarily “good” either for individuals or for societies. These theories suggest a form of Darwinian selection at the level of values: values that “persist” and that are culturally-shared are values that manage to be transferred and exchanged. To that extent, values supported by culture can lead to behaviors that have no or limited benefits, both from an individual and collective perspective (e.g. value of “self-sacrifice”, proselytism). This mechanism is facilitated by the fact that values supported by culture are taught when individuals are young and hardly change afterwards because they are hard to reconsider rationally (high complexity for doing it and not complying with cultural values generally leads to punishment).

Nevertheless, cultural values generally support sustainable individual and collective outcomes. Indeed, values supporting better individual and collective welfare tend to be learned more easily. Furthermore, cultural values should influence societies such that these societies are successful enough for maintaining and transmitting values (e.g. educative system). This means that not all cultural values are appropriate for supporting coordination, but most values are and can be used as examples.

Summary

This section introduced a core mechanism of the influence of culture on coordination: the value-system. This system is both sensitive to culture and it particularly impacts individual decisions in situations of coordination and in a wide range of environments.

4.2 Modeling Value-Systems

Section 4.1 introduces a key mechanism of the influence of culture on coordination: value-systems.

The current section investigates how to model value-systems and the influence of culture on them. In order to do so, we propose an architecture for modeling value-systems and how to relate them to culture.

In order to build this architecture, we should first determine what we want to model. Indeed, as presented in Section 4.1, the influence of value-systems on decisions is particularly complex. Nevertheless, since we only want to replicate the properties the influence of values that matters for achieving coordination, we do not need to capture all this complexity for achieving our goal.

We aim at using the influence of culture on value-systems for supporting coordination. As a parallel, we want to use value-systems as a support for designing agents, like norms

and organizations help to design agents capable of achieving good coordination. In our case, value-systems should drive agents towards a “good” direction (according to values) when making decisions.

In order to achieve this aim, we propose a model of value-systems that helps making *value-compliant decisions*. This model aims at determining whether a decision is in line with a value system or not, while skipping the details of the influence of value-systems on cognitive processes.

We acknowledge that this model skips the rich but complex deep cognitive features of values proposed by Miceli and Castelfranchi (1989) (e.g. values taking over rational processes while evaluating). Nevertheless, this level of detail is out of the scope of this thesis. We only need to replicate the influence of culture at a behavioral level. Our model is sufficient for replicating the influence of culture on coordination. Furthermore, this model can be more easily embedded within standard decision processes, as we do in Section 4.3.

van der Weide (2011) provides a relevant architecture for modeling value-systems which can be used for making value-compliant decisions. This architecture is relatively complete and our architecture is broadly inspired by this work. We expand this architecture towards making decisions.

In this section, we first introduce an architecture for modeling values in Section 4.2.1. Based on this, we introduce an architecture for modeling value-systems, in Section 4.2.1. Finally, we relate this model to the influence of culture in Section 4.2.2.

4.2.1 Modeling Values

As indicated in Section 4.1.2, in essence, values indicate “good” and “bad” features of things (e.g. object, idea, decision) in general (e.g. honesty, obedience, discipline, autonomy, strength, robustness). These features can then be used within deliberations for deciding whether a thing is “good” or “bad” (e.g. this item is “good” because it is robust). These features can also be used as a point of comparison: the thing t_1 is better than t_2 from the perspective of value v .

The set of values is formally represented by V . Each value $v \in V$ is related to a set of possible evaluations represented by $eval_v$. This set indicates the range of possible evaluations for this value (e.g. “{very bad; bad;neutral;good;very good}”).

The set of evaluations is internally ordered (e.g. “very good” is better than “neutral”). This order is formally represented by the total order $<_v$ over $eval_v$. If $e_1 <_v e_2$ then e_2 is “better” than e_1 according to v .

Values evaluate things. Formally, we represent that with the function $evaluate_v : T \rightarrow eval_v$ where T informally represents a set of “things” (e.g. alternatives of a decision, ideas, items).

Values can compare things (assuming that they are comparable, like comparing alternatives of a given decision). In considering $t_1 \in T$ and $t_2 \in T$ as two things, then t_2 is better than t_1 according to $v \in V$ if

$$evaluate_v(t_1) <_v evaluate_v(t_2)$$

Value-Systems

Value-systems consist in a combination of multiple values. In particular, value-system should indicate the relative importance between values.

Formally, we represent value-systems by a set *eval* of evaluations, the partial order $<$ which orders *eval* and the function *evaluate*: $T \rightarrow eval$, which evaluates the “goodness” of a thing.

The set of *eval* consists of all possible tuples of evaluations for each value. Formally, if $V = \{v_1, \dots, v_n\}$ then *eval* is:

$$eval = eval_{v_1} \times eval_{v_2} \times \dots \times eval_{v_n}$$

The evaluation function *evaluate* combines the evaluations of all values. This function is represented by:

$$evaluate(t) = (evaluate_{v_1}(t), \dots, evaluate_{v_n}(t))$$

Finally $<$ orders possible evaluations. This function combines multiple preference functions as depicted by Brafman and Domshlak (2013). In particular, the property known as the “*ceteris paribus*” applies. This property is informally introduced follows. For any pair of things $t_1 \in T$ and $t_2 \in T$, if t_1 is better evaluated than t_2 for a given value and is either better or incomparable according to all other values, then t_1 is better than t_2 according to $<$.

Based on this, the task from a design perspective consists in completing $<$ (e.g. what is the best between (v_1 : “very good”, v_2 : “neutral”) or (v_1 : “good”, v_2 : “good”)).

This model of value-systems aims at being simple and sufficient for performing value-compliant decisions. Nevertheless, this simple representation can be further expanded. For instance, van der Weide (2011) suggest using specification trees for relating abstract values to more concrete values (e.g. relating obedience to compliance and information reporting). This technique is further described in Section 7.2.4 which introduces techniques for efficiently designing cultural values for supporting coordination.

General Guideline for Using Value-Systems for Making Decisions

Our model of value-systems can be directly used for making value-compliant decisions. We only assume that, for a given decision, agents can determine a set of rationally-equivalent *alternatives*. This set of alternatives corresponds to multiple possible outcomes that are difficult to handle rationally (e.g. too expensive, too similar alternatives): no alternative is clearly better or worse than the others. This situation can occur often in the type of environment we investigate, due to complexity or uncertainty or because of interactions.

Our model can then be used for making value-compliant decisions. Formally, given a set of rationally-equivalent alternatives $Alt \subset T$ for a decision, the alternative that is the most value-compliant is:

$$\operatorname{argmax}_{alt \in Alt} evaluate_d(alt)$$

where *max* is defined based on $<$.

For instance, consider two plans: finishing work as early as possible to get back home versus taking more time for improving quality. Both alternatives are rationally acceptable (both match agent goals and no one dominates the other). In that case, alternatives can be compared by using the value-system of the individual for making decisions. For instance, since the value of “professionalism” is more important than “being relaxed”, then the plan “taking more time” is preferred over the other.

4.2.2 Modeling the Influence of Culture on Value-Systems

Value-systems constitute an entry-point for culture. Cultures influence value-systems. Theories indicate multiple properties of the influence of culture on value-systems that are applicable here.

Culture is known to influence the range of values adopted by individuals and their relative importance. In general, culture supports values that are oriented towards social behavior. Culture should impact without completely determining the value-system of individuals. Individuals expect their value-system to be shared. Likewise, individuals can have different ways for evaluating a given value, even if individuals tend to evaluate the same value similarly.

In our architecture, a way of representing the influence of culture consists in having all agents sharing the set V of values, $<$ and for all values $v \in V$ $eval_v, <_v$. About the evaluation functions $evaluate_v$, we only assume that these functions remain conceptually coherent with the meaning of v . Nevertheless, we leave free the implementation of $evaluate_v$ which is dependent on the decision process of the agent.

Summary

This section introduces an architecture for modeling values and value-systems in order to perform value-compliant decisions. In addition, we relate the influence of culture on models generated by this architecture.

4.3 Modeling Culturally-Sensitive Agent Decision Processes

This section proposes an architecture for modeling *culturally-sensitive agents* capable of making decisions in *situation of coordination*. This architecture is to be used in Chapter 7 as an example for showing that the influence of culture on coordination can be replicated within MASs.

Our architecture requires two features: sensitiveness to culture and capability for making decisions in situations of coordination. In order to integrate the first feature within our architecture, we integrate the influence of value-systems, introduced in Section 4.2, on decisions. For the second feature, we rely on a traditional decision architecture for making decisions in a situation of coordination: the BDI decision architecture presented in Rao and Georgeff (1995), introduced below. Therefore, we enrich the traditional BDI architecture with value-systems.

Decision models based on the BDI architecture are primarily useful for agents embedded in an environment and making successive decisions leading to actions. When a decision occurs, the agent receives a set of percepts, updates its memory and determines which action to perform. The BDI architecture particularly aims at performing “practical reasoning”. Practical reasoning consists in setting a goal (leading the environment to a given set of states) and in making decisions with the aim of achieving this goal. In order to achieve practical reasoning, the BDI architecture consists of two parts: a memory structure and a decision loop. The *memory structure* orders agent memory in three structures that support practical reasoning: Beliefs (information about the world), Desires (the goal to achieve) and Intentions (what to perform next). The *decision loop* consists, when facing a decision to be made, in performing a sequence of internal decisions for updating the memory structure, particularly for building plans for achieving the goal.

Why relying on the BDI architecture? The BDI decision architecture is one of the most appropriate available models for making decisions in a situation of coordination, particularly in complex environments. Furthermore the BDI model has been explored for a relatively long time and is relatively broadly used. The BDI model matches conceptually relatively well with values. Finally, values complements relatively well the support provided by the BDI architecture for making decisions. Indeed, the BDI architecture offers a strong goal-directed rational basis for making decisions that benefits from the heuristic efficient yet irrational support provided by values when rational limits are reached.

In this section, we present the two parts of this architecture in succession. The memory model is introduced in Section 4.3.1. The decision loop is introduced in Section 4.3.2.

4.3.1 Memory Model

The memory model represents the internal structure of the memory. In the traditional BDI model, this structure is composed of three bases: beliefs, desires and intentions. In this model, we add one base: the value-system.

Following subsections introduce each base independently.

Beliefs

Beliefs represent what is believed to be true (e.g. the block *B* is on top of the block *C*). In general, beliefs rely on a model of the environment. They are used for estimating the state in which the world is found currently. Beliefs are updated in the “process event” phase from the update loop, when new percepts are received. In practice, beliefs can be modeled by using first order logic as in Dastani (2008).

Beliefs are not necessarily exact. Indeed, an agent may believe something false due to misperception, misinterpretation or an inappropriate model of the world.

Desires

Desires represent what the agent wants to pursue. In theory, desires are abstract drives. In more practical implementations of the desire base, desires are concrete states or properties about the environment that the agent wants to achieve.

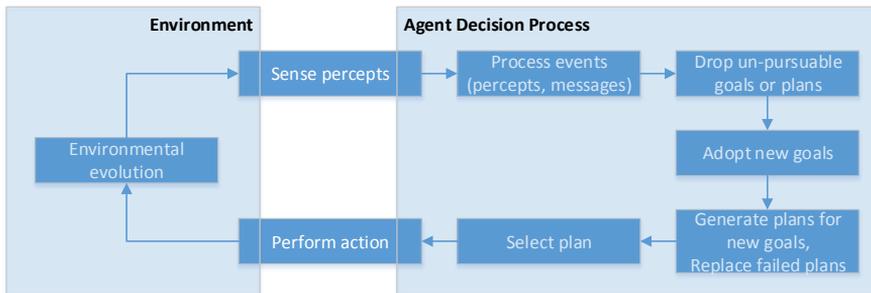


Figure 4.2 – *Decision Loop of our BDI Architecture*

Desires are related to beliefs. A desire which is believed to be fulfilled should be dropped. Likewise, desires that are believed to be unfeasible should be dropped. A step of the update loop is dedicated to dropping goals.

Agents can also adopt new goals or revise a goal they are currently aiming for. These goals are driven by the general desires of the agent. Certain properties between adopted goals should hold, such as goals should not contradict each other.

Intentions

Intentions represent courses of actions to be pursued and committed to in order to achieve goals.

In practical settings, intentions are represented by plans. Plans are represented by algorithms that determine which actions to perform, with the possibility of branching and looping depending on beliefs. Plans are generated by the decision loop. Each plan is driven towards achieving a given goal.

The decision process tends to maintain a coherent plan base, for instance by dropping plans when the related goal is dropped, by making sure that conditions required for performing a plan are still met, by making sure that plans do not contradict each other.

Value-Systems

As depicted in Section 4.2.1, value-systems represent the set of values adopted by the agents and the relative importance between these values. Value-systems can in particular be compared for supporting decisions, when multiple rationally-equivalent alternatives are available.

4.3.2 Decision Loop

The decision loop, represented in Figure 4.2, is the sequence of internal decisions that the agent performs for deciding which action to pick. The decision loop presented in this section is the standard BDI decision architecture. Value-systems influence this decision loop in influencing the resolution of internal decisions. An example of a generic mechanism

for integrating the influence of value-systems on decisions is introduced in Section 4.2.2. Following subsections introduce decisions that are made at each step of this loop and how value-systems influence this decision.

Process Events (Percepts, Messages)

The first step of the decision loop consists in processing incoming percepts and events. The aim of this step consists in updating the belief base with recent observations. These updates can alter other bases at later steps of the decision loop (e.g. receiving a message giving a goal, finding out that a goal cannot be achieved).

This step raises decisions to be made, particularly in practical settings that require to manage the size, the efficiency and the consistency of the belief base. First, agents have to decide which beliefs they store. For instance, they may not want to keep the complete camera record, but they may want to store visual information if they observe a crime. Second, agents can revise beliefs with new clues (e.g. this door is believed to be closed, what if a new record indicates it is open?). Third, agents can infer beliefs from raw data, which are more convenient to process (e.g. replacing “seen footprints” by “someone was here”).

Influence of Value-Systems: Value-systems can influence these decisions. Value-systems influence the importance given to beliefs and thus the decision to store them (e.g. behaving according to protocol is important, thus the agent should remember whether exceptions occurred when acting). Value-systems can influence belief revision (e.g. uncertainty is very bad, better get more evidence before changing beliefs). Value-systems can also influence inference mechanisms (e.g. if tradition is important, the message “task processed” can be translated into “task processed exactly according to standards”).

Drop Goals or Plans

Adopted plans and goals can be dropped. In general, goals and plans are dropped when they are found to be impossible to achieve, to be already achieved or when the goal attached to a plan is dropped.

This can also happen when the desire evolved and existing goals or plans are no longer in line with new desires (e.g. “replacing the tape of the typing machine” is no longer relevant if the agent is provided with a computer and a printer). This can also happen when opportunities are available (e.g. “get a cup of coffee” can be dropped if promised a lot of money for a day without coffee). Likewise, goals and plans can be dropped if more desirable goals and plans are available or if they contradict with other goals and plans.

Influence of Value-Systems: Dropping goals and plans raises numerous decisions that are sensitive to values. Values influence the commitment to plans and goals in general (e.g. LTO-CD-related values). In addition, value-systems deeply influence the importance (or the “goodness”) of a plan or a goal and thus the decision for dropping it.

Adopt New Goals

Agents can adopt new goals. Goals are mostly aimed at satisfying the desires of the agent but can also be adopted as sub-goals for achieving another goal.

Adopting new goals raises numerous decisions. New goals should satisfy at best desires, not contradict each other, and match the environment. These constraints can raise

complex decisions for determining how far a goal satisfies desires, which goals to pursue and which to drop.

Influence of Value-Systems: Value-systems influence this decision by influencing the evaluation of the importance of a goal, but also by considering the challenge level of a goal (e.g. risk vs. rewards).

Generate Plans

Plans are generated such that all goals are related to a plan. New plans can be generated when new goals are adopted or when an existing plan is dropped (e.g. because this plan has failed).

Generating new plans raises several decisions. New plans have to match existing goals and other plans. Furthermore, several plans can be used for achieving the same goal, raising numerous alternatives. Sometimes, adopting a plan can force to revise other plans.

Influence of Value-Systems: As for goals, value-systems can influence the selection when multiple alternative (sets of) plans can be adopted for achieving (sets of) goals. For instance UAI-CD-related values support the selection of plans based on protocols or not, PDI-CD-related values support higher compliance with instructions provided by leaders, possibly requesting to be given plans by leaders.

Select Plan

Agents may have a set of plans to perform at a given time. In this case, the agent has to select a plan to act upon.

This selection involves again numerous decisions. The various plans can fulfill goals of various importances. The current situation can raise opportunities for triggering several plans.

Influence of Value-Systems: As introduced in previous steps of the decision loop, value-systems influence the selection of plans when multiple alternatives are available, in considering the inadequacies of plans with values and the importance of related goals (e.g. focus on plan p because this plan achieves the goal set by my leader and obedience is important).

In addition, value-systems influence the strategy for achieving plans (e.g. focus on one plan at a time for completing it quickly vs. maximize opportunity taking, possibly delaying the completion of plans).

Perform Action

This step consists in performing the first action of the selected plan. Since there is only one plan selected no decisions are to be made. As a consequence, value-systems do not impact on this step of the decision loop.

Conclusion

This section presents an architecture for designing culturally-sensitive agents capable of making decisions in situation of coordination. This architecture expands the BDI archi-

ture, which is adequate for performing practical reasoning, with value-systems, which provide an entry-point for integrating the influence of culture.

Value-systems and the BDI architecture appear to match relatively well. On the one hand, the BDI architecture provides rich “rational” support for making goal-directed decisions. On the other hand, value-systems provide rich “irrational” heuristics for making decisions when outcomes of these decisions become too complex. These heuristics are particularly relevant in our case because we aim at tackling complex environments. Thus, values provide irrational yet relatively accurate indications, which complement the rational but support provided by rational decision processes, which is limited by computational complexity.

4.4 Culturally-Sensitive Decision Aspects

In Section 4.3, we introduce a general architecture for making culturally-sensitive decisions. On top of this architecture, numerous aspects can influence decisions, particularly in situation of coordination (e.g. uncertainty, relationships between agents). These aspects are also known to be impacted by culture. In order to illustrate that our architecture can be used for replicating the influence of culture on coordination, we show that our architecture can replicate these cultural phenomena.

To that extent, the current section introduces the main decision aspects that we found in our studies to be influential in situation of coordination and that are known to be particularly sensitive to culture. During our studies, we found out four core decision aspects that match this description: identity/relationships, norms, uncertainty and learning. These aspects are the *core* ones, which are particularly considered for decisions in situation of coordination. Nevertheless, we acknowledge that more decision aspects, less often used, can also be considered (e.g. motives, as presented in Chapter 5).

In this section, we further describe each aspect, we show its impact in situation of coordination and how culture influences it. Finally, we introduce how this aspect can be implemented within our architecture.

Section 4.4.1 introduces identity and social relationships as a decision aspect, Section 4.4.2 introduces norm as a decision aspect, Section 4.4.3 introduces uncertainty as a decision aspect and finally Section 4.4.4 introduces learning as a decision aspect.

4.4.1 Identity and Relationships

Agents can be considered in considering their *identity* (e.g. expertise domain, performance, role), which indicate the core features of an agent. In addition, agents can create special *relationships* between each other (leader-subordinate links, workflow links). Both decision aspects tend to influence dramatically how agents interact with each other (e.g. responsibilities, what to expect, importance).

Identity and relationships influence coordination. They are particularly used when coordinating with organizations, as introduced in Section 2.1.2: identity is formalized through roles and relationships through formal structures in formal organizations and both aspects are further enriched by informal organizations (e.g. status, reputation, friendship).

From the perspective of individual decisions, identity and relationships deeply influence communication. Depending on their identity and their relationships with other agents, agents can issue different messages (e.g. orders to subordinates, suggestions to colleagues), raise different communication patterns (e.g. share everything with everyone in a group, specific communications) and process messages differently (be more ready to comply with instruction of a leaders, be more cautious with cooperators).

Identity and relationships are related to culture. Their definition partly depends on culture (e.g. leader-subordinate relationships tend to be a preferred form of interaction in high PDI-CD organizations as introduced in Figure 2.3). Furthermore, the interpretation of these relationships is particularly culturally-dependent. For instance, “being a leader” in high PDI-CD organizations is tightly related to being the one in charge of making all decisions. In low PDI-CD, being a leader is more related to being in charge of connecting subordinates with each other and let them decide how to solve tasks.

Theory

Through local interactions and coordination, individuals can be related to an identity and create relationships with each other.

Four main culturally-sensitive and coordination-influenced aspects of identity/relationship can be considered, based on theories such as Hofstede et al. (2010a): *leadership* (presence of a leader-subordinate link), *membership* (agents being related to similar groups, e.g. subordinates coordinated by the same leader), *expertise* (accuracy when making decisions for a given decision range) and *good performance* (e.g. concrete acknowledged performance of the agent). These are the main aspects we considered. More can be added to that (e.g. friendship as proposed by Mc Breen et al. (2011)), but these variables do not appear to influence coordination as well.

These aspects influence the value given to a relatively general aspect: *status* in Section 2.1.2. Status influences the amount of concessions that an individual is willing to make towards another individual. Individuals with higher status tend to be given more importance when taking decisions and their information is better trusted (e.g. subordinates are more compliant to orders of leaders with high status).

Status impacts how to influence others and how to react to influences from others (e.g. subordinates are expected to be more compliant with instructions provided by their leaders). This influence can be conditional to certain situations or messages (e.g. the additional importance given to communication from experts is relatively bounded to their expertise domain).

Culture impacts the influence of aspects of identity/relationship on status. For instance, PDI-CD influences the importance given to the leader-subordinate relationship on status. The same applies to experts with the UAI-CD, good performers with the MAS-CD and group members with the IDV-CD.

Integration within Our Decision Architecture

Within our decision architecture, we model aspects of identity/relationships (leadership, membership, expertise, good performance) within the belief base. These aspects can be represented with a network that instantiates four social variables: subordination, expertise, performance and membership. These variables are combined together and weighted

by culture (e.g. importance of values of obedience, achievement, tradition, respect of social order) for providing the value of an additional variable: status.

Processing Events: Acknowledged influences (e.g. messages) can generate the decision to adopt new beliefs and drop or revise current beliefs.

In this case, the status of the sender influences whether this message is accepted as a belief or takes over existing beliefs. For instance, if a_1 said “the door is blue” and a_2 said “the door is red”, the agent can prefer to believe a_1 which has higher status. Culture can indirectly impact such a decision. For instance, a_1 is the leader and a_2 is a good performer. The opinion of a_1 can be preferred if obedience is culturally important while the opinion of a_2 can be preferred if achievement is culturally important.

Percepts can update information about the social identity of other agents: changing role (e.g. promotion), updating performance (e.g. new information about successes or failures of someone), acknowledging expertise (e.g. discovering a new field of expertise) and changing groups (e.g. interacting more with another group).

Dropping Goals and Plans: Messages can suggest dropping a goal or a plan or proposing alternative goals or plans that would replace the current one. The decision to drop a goal or a plan is sensitive to the status of the sender in general and to its identity. Agents should give more importance to instructions provided by leaders, even if leaders are given relatively low status. Nevertheless, suggestions from agents with higher status can be adopted more easily as goals. As an illustration, agents may drop goals given by their leaders if an agent with high status tells them that it is not feasible, depending on culture (e.g. more or less easily depending on the importance given to obedience).

Adopting Goals: As for dropping goals and plans, agents are more likely to adopt goals that are suggested by agents with high status, particularly if these agents are leaders.

Generating Plans: As for adopting goals, agents prefer to have plans that match those proposed by other agents, depending on the agent's status, but also its expertise if relevant and the presence of subordination links. Furthermore, agents are more likely to give advice or instructions to agents with lower status.

Selecting Plans: Messages can influence the relative importance given to a goal or to a plan and thus whether this plan is to be triggered. These messages are given more or less credit depending on the status of the sender.

Summary

In this section, we introduced basic aspects and a model of the identity/relationships decision aspect. This aspect is crucial for supporting coordination, in providing a richer social context for supporting interaction-related decisions. Furthermore, this aspect is known to be particularly culturally-sensitive.

In taking this aspect in consideration, we can understand better why a given organizational structure leads to a possibly wide range of collective outcomes due to different cultural backgrounds. A given organizational structure supports decisions that are handled differently by different cultures. For instance, the dilemma between “trying to fix locally a problem or report it” is solved differently depending on the way leaders and subordinates consider each other. In high PDI-CD culture, subordinates are more likely to give too high importance to leaders and not bother them with details, thus fixing locally. In low PDI-CD culture, a reverse decision is likely to be made.

Parts of this decision aspect are used in this thesis, for performing simulations. In Chapter 5, this decision aspect highlights the culturally-sensitive importance of “expertise” and “leadership” on interactions. In Chapter 6, this decision aspect also highlights the “leadership” social aspect.

4.4.2 Norms

Norms constitute a second group of culturally-sensitive and coordination-influencing decision aspects. Norms, introduced in Section 2.1.4, indicate behaviors to be complied with. Norms also induce norm-related decisions (e.g. violating and enforcing norms).

Norms influence coordination and are often used as a coordination technique. Norms aim at directly enforcing behaviors support coordination (e.g. cross when the light is green) or preventing coordination-damaging behaviors (e.g. do not cross when the light is red). Additionally, by determining the behaviors of agents, norms can be used for creating expectations about other agents and can thus react accordingly.

From the perspective of culture, norms are known to be culturally-sensitive Mc Breen et al. (2011). Norms support individual and collective decisions that are in line with cultural values (e.g. rules of fair competition in high MAS-CD culture). Furthermore, values are known to support norm compliance and norm enforcement of norms that are in line with culture (e.g. “do not cross the line because this is not fair”).

Theory

Culture, through values, influences the decisions related to adoption, compliance and enforcement of norms. This influence can be expressed in relation with the type of norms and with the content of the norm.

Types of Norms: Two main types of norms have been acknowledged in literature: legal and social norms. *Legal* norms can be seen as regulations, at restricting agent behavior, generally, for supporting coordination. *Social norms* correspond to the acknowledgment of “regular” behavioral patterns. These two types of norms share similarities with regard to their influence on decisions (e.g. create expectations, compliance is desirable) but rely on different social mechanisms (e.g. how is the norm determined? How is it enforced?).

Values influence the relative preference between the two types of norms. As highlighted in Mc Breen et al. (2011) and by the UAI-CD, values (e.g. conformism, compliance) influence the relative importance given to legal and social norms and thus the compliance with these norms.

Content of the Norm: Norms can be related to a “content”, which indicates the restriction conveyed by the norm. This content can be related to underlying reasons for complying with this norm (e.g. the norm “pay for the bus” exists for maintaining bus services or for paying bus drivers). These reasons can then be evaluated according to values (e.g. “paying bus drivers” is good according to equity).

Furthermore, the content of the norm can indicate how a norm can be controlled and which sanctions can happen if a violation occurs. Again, values can evaluate sanctions (e.g. being caught implies paying a fine, get off the bus and public shame). This way, culture can influence norm-compliance.

Integration within Our Decision Architecture

Modeling the influence of norms on decisions is a broad and complex topic. For instance, the decision architecture proposed by Dignum et al. (2001) requires additional modules.

To that extent, in the current section, we only briefly describe possible connections between values, norms and decisions, focusing on core aspects. Nevertheless, we acknowledge that the connection between these aspects is conceptually rich as theorized by Miceli and Castelfranchi (1989) (e.g. using values for motivating the compliance with a norm). We leave this rich ground for research as will be indicated in Section 8.3.3

We can expand our model in relating agents to a “norm-base”. This base contains the set of externally-given legal norms and social norms that are either given at the start or inferred by the agent. Norms are related to the behaviors they forbid, the reason of their existence and consequences in case of violation. Furthermore, agents are assumed to be capable of estimating whether a plan or a goal can lead to norm violation.

Process Events: Perception is used for determining whether the agent violated a norm or whether another agent did so. Likewise, new social norms can be inferred from perception.

Drop Goals or Plans: Norms influence how agents reconsider their plans or goals. The main influence happens when a plan or a goal can lead to norm violation.

Culture, through values, influences the reasons for caring or not about norm violation. If the content of the norm is streamlined with values, then values should support the decision to drop the norm-violating goal or plan. Likewise when the consequences of the expected violation are very “bad” according to values.

Adopt new goals: When adopting a goal, the agent evaluates whether this goal leads to norm violation. If so, the agent evaluates whether pursuing this goal is more important than complying with the norm in considering the relative importance given to the norm and to the consequences of its violation.

Values influence the relative importance given to the norm and its violation. Thus, values impact on the threshold at which the agent accepts to select norm-violating goals.

Generate New Plans: Norm violation is a negative feature of a plan (unless the agent purposefully wants to violate the norm). Thus, agents generally prefer plans that avoid norm violation. Nevertheless, depending on the goal to be fulfilled agents can decide to perform a plan that entails norm violation.

As for goals, values can indirectly influence this decision by influencing the relative importance given to norm compliance (depending on the match between the norm and values).

Select Plans: Plan selection also takes into consideration whether plans violate norms and possible consequences of this violation.

As for plan generation and goal adoption, values alter the relative importance given to norm violation.

Conclusion

This section relates norm-related decisions to the influence of culture through values. Values influence in particular decisions related to norm violation, norm compliance and norm enforcement. As we studied in Vanhée et al. (2011), norm-related reasoning is relatively challenging to implement within a rational BDI agents (e.g. hard to predict indi-

vidual and collective consequences). To this extent, value-systems support decisions by providing principles for determining how “bad” a violation is and how “important” compliance is.

4.4.3 Uncertainty

Uncertainty is a decision aspect that is at play when consequences of decisions are hard to predict. Multiple reasons can cause uncertainty, for instance inherent randomness, computational complexity, an incomplete model for predicting consequences, partial observability or the influence of other agents.

Uncertainty is frequent in situations of coordination, particularly when the environment is complex or dynamic. Many techniques for achieving coordination aim at re-introducing a certain degree certainty within the system (e.g. the contract-net protocol helps determining what to achieve and what to expect in return). Some techniques aim at perfectly handling uncertainty (e.g. DEC-POMDPs), but these techniques have too high requirements for being implemented in our case (e.g. a complete model of environmental dynamics, high complexity).

Uncertainty is also a core aspect of culture as highlighted by the UAI-CD. This dimension highlights the influence of culture on the sensitiveness to uncertainty and culturally-sensitive techniques that are deployed for handling uncertainty (e.g. using rules, very detailed communication, trust in experts). To some extent, the MAS-CD is also related to uncertainty. In low MAS-CD, individuals tend to avoid uncertainty that can lead to failure in avoiding conflicts. Conversely, high MAS-CD supports higher risks for higher rewards reasoning.

Theory

Uncertainty arises from multiple sources. Uncertainty can be raised by assumed randomness (e.g. the result of a dice roll impacts on the decision), because predictions are too complex to make (e.g. expected evolution of the stock market), because other agents can influence the outcome of the decision (e.g. coordination games) or because the current situation is only partially known (e.g. the door may be opened or not).

Uncertainty influences decisions in many respects. In order to avoid exploring aspects irrelevant with regard to this thesis, we focus on aspects of uncertainty that are sensitive to culture. To our knowledge, culture influences aspects of uncertainty caused by partial observation, uncertain consequences of actions and collective action. These aspects of decisions can be handled by performing actions that provide information about the environment (e.g. checking), provide more certainty about consequences of actions (e.g. relying on plans that are well-grounded) and about others (e.g. knowing about the plans of others and making sure they comply with their plan).

Culture can be related to two tendencies for handling uncertainty when making decisions, particularly in the context of coordination. A first tendency related to high UAI-CD consists in assuming and maintaining low uncertainty. Maintaining low uncertainty can be achieved in relying on standards and extensive communication. A second tendency consists in making little assumptions and letting individuals re-creating locally some certainty (low UAI-CD).

Multiple more concrete tactics exist for implementing these two tendencies. In the first tendency, agents can maintain their relative certainty by following protocols or standards. Standards are generally based on strong assumptions about the environment and lead to stable and expectable outcomes. Thus, standards satisfy all sources of uncertainty: agents have more certainty about the current state of the environment, about consequences of their decisions and about decisions of others. When unexpected events occur, agents tend to get back to a standard situation or warn other agents about these irregularities. In the second tendency, agents tend to locally get sufficient information for achieving their tasks. They tend to rely on more adaptive plans for achieving their goals, increasing the range of possible outcomes.

Culture also influences the contextual importance of uncertainty. For instance, in high PDI-CD organization, managing uncertainty related to satisfaction of the supervisor is given more importance than in low PDI-CD organizations. Likewise, in low MAS-CD organizations, individuals prefer to avoid uncertainty about failures.

Integration within Our Decision Architecture

Culturally-sensitive influences of uncertainty on decisions can be integrated in multiple places within our architecture.

Process Events: In complex and dynamic environments, observations and messages may be insufficient for exactly determining in which situation the agent is. Nevertheless, in order to operate efficiently, agents may create inferences based on these partial observations (e.g. if my light is green, I assume that the light of other lanes is red and thus other cars will not cross).

Uncertainty can influence the decision to create inferences or not. Depending on culture (e.g. traditionalism, security), agents are more or less prone to assume that the environment is standard by default or not.

Drop Goals or Plans: Uncertainty can raise the decision to drop a goal or a plan. Agents may be more or less sensitive to the uncertainty generated by a goal or a plan. In this situation, such a goal or plan can be dropped for selecting an alternative which provides more certainty with regard to its achievement and for others (e.g. because more standard).

Culture, through values (e.g. taking risks is “bad”, normal situation is “good”), influences the sensitivity of agents to drop uncertain goals and plans.

Adopt New Goals: Uncertainty can influence the adoption of goals. The certainty to achieve a goal is a desirable quality of a goal, which may be traded off for satisfaction (e.g. more certainty for a bit less of satisfaction). Values can influence this tradeoff (e.g. MAS-CD-related values supporting assertive “high risk high reward” decisions or modesty). Likewise, culture can support the decision to pursue standard or “traditional” goals or explore a new type of goals.

Generate Plans: The influence of uncertainty on plans is similar as for goals. Agents can be more or less likely to adopt risky or standard plans according to their values for lowering their own uncertainty.

Conclusion

This section shows how uncertainty can influence decisions related to coordination and how culture influences the importance of uncertainty. This section shows how values solve

several social dilemmas (e.g. maintaining low uncertainty anytime vs. locally re-creating certainty). This decision aspect is implemented within a simulation in Chapter 5 in considering the influence of the MAS-CD on uncertainty management and thus on coordination.

4.4.4 Learning

The last decision aspect that we investigate is learning. Learning consists in adapting decision processes to the environment. In practice, learning is generally directed towards improving agent performance in altering internal decision variables.

Learning is a coordination technique as introduced in Section 2.1.4. In this case, agents generally aim at altering their internal variables for optimizing a collective performance criterion.

Learning is known to be culturally-sensitive. First, the performance criteria to optimize are culturally-driven (e.g. having good individual performance in high MAS-CD organizations versus being a good team-worker in low MAS-CD organizations). Second, the way learning is performed and learning outcomes are also sensitive to culture (e.g. quickly learning of a few useful rules vs. slowly learning over time, as suggested by the LTO-CD).

Theory

As presented in Chapter 3, culture influences *how* individuals learn. Low LTO-CD values (e.g. being a quick learner is “good”) support teaching and learning simple rules, which are easy to learn and to apply. Conversely, high LTO-CD values (e.g. dedication, commitment) support more “intuitive” learning based on numerous trial-and-errors, requiring more effort to learn, are harder to transfer but are more rewarding in the long run.

In addition, learning is based on a reward mechanism, which is sensitive to values and thus to culture. Indeed, individuals tend to be rewarded (intrinsically or not) when performing “good” behaviors, which is driven by values (e.g. individual performance, team play, match standard, in due time).

Integration within Our Decision Architecture

These forms of learning can be integrated within our model through two learning modules, as part of the belief base. The first module is a mechanism for performing quick learning (e.g. trees, logical inference). The second module is a mechanism requiring more data but leading to higher learning quality (e.g. neural networks). These modules are then applied for making decisions, for learning about the world model (e.g. this route is overloaded between 16:00 and 18:00) or about plans (e.g. the plan “lifting the table alone” does not provide good results).

Values influence the learning criteria. Values subjectively indicate how “good” is an outcome and thus how to integrate this learning experience.

These mechanisms are updated during the *process event* phase, by checking whether expectations about the world model are correct and by updating plans and goals depending on their achievement or not. Learning, by updating the world-model can influence all possible decisions, particularly deciding to drop a goal and to select goals and plans to adopt.

Values influence decisions that support a form of learning over the other. High LTO-CD related values (e.g. commitment, dedication) support decisions that lead to repetitive trials (e.g. prefer adopting similar goals), supporting the use of the long-term learning module. Low LTO-CD related values stress less on repetitive behavior, supporting the learning of simpler rules.

Summary

This section shows evidence that our architecture can be used for replicating the influence of culture on coordination. We have indicated for four decisions aspects that are both culturally-sensitive and which impact coordination how these aspects can be introduced within our model.

4.5 Chapter Summary

This chapter performed the first two steps towards highlighting that the core properties of the influence of culture on coordination can be replicated within MASs.

In the first step, we highlighted a central mechanism of the influence of culture on coordination in human societies: value-systems. Culture influences value-systems of individuals, by promoting a set of common values (indicating what an individual considers as “good”) and giving them a relatively similar importance. Then, value-systems influence many decisions performed by individuals. This influence is particularly visible in situation of coordination (e.g. promoting some norms, creating expectations about others, agreement about what is important).

In the second step, we integrated this mechanism within MASs by modeling it. We first introduced an architecture for modeling value-systems, the entry-point for culture. Then, we integrated these value-systems within an agent decision architecture capable of making decisions in situations of coordination. Furthermore, we have introduced four additional decision aspects (identity/relationships, norms, uncertainty, and learning) that are culturally-sensitive, coordination-impacting and can be captured by our decision architecture. Together, this decision architecture provides a core entry-point for integrating the influence of culture on coordination within MASs.

The next step, which will be studied in Chapter 5 and Chapter 6 consists of validating this model. These chapters show that our architectures can build models that empirically replicate the core properties of the influence of culture on coordination, at multiple levels (individual decisions, collective outcomes). This model will be used for building simulations in order to better explain the individual-to-collective link of the influence of culture on coordination in Chapter 6. Finally, this model of the influence of culture on coordination illustrates that such a model can be done, while highlighting the core underlying mechanisms at play. Based on this, we will study in Chapter 7 how to use this model and these mechanisms in order to coordinate agent societies.

Answering Research Question 2.a: *What Are the Key Mechanisms of the Influence of Culture on Coordination?*

This chapter introduces a core mechanism of the influence of culture on coordination: value-systems.

Culture indirectly influences individual decisions. Culture influences value-systems by raising a set of cultural values and by influencing the relative importance given to values. Values heuristically indicate what is “good” and “bad” in considering certain general features (e.g. honesty is “good”). This way, value-systems influence a wide range of decisions by suggesting alternatives of this decisions that display the “best” combination of features.

Value-systems supported by culture are particularly influential and supportive for coordination. First, because supported values are related to social behaviors. Second, because they are shared and expected to be shared. Thus agents can create strong expectations decisions of other agents (e.g. not performing “bad” behaviors) and are driven towards considering similar outcomes (e.g. agree to ask and reveal private information because uncertainty is collectively agreed-upon to be “very bad”). Third, because they can be relatively abstract, thus being applicable in a wide range of situations. Fourth because they provide a common basis for creating more concrete coordination techniques (e.g. value-supported norms).

Answering Research Question 2.b: *How Can We Build Practical Models of the Key Mechanisms of the Influence of Culture on Coordination?*

This chapter proposes a model for replicating key mechanisms of the influence of culture on coordination. More precisely, this chapter proposes an architecture for modeling value-systems, which are culturally-sensitive. Then, this chapter proposes an agent-based decision architecture that both integrates the influence of value-systems and can be used for making decisions in a situation of coordination.

Our architecture for building value-systems consists of two parts. The first part is, a set of values. Each value can evaluate whether ‘things’ (in our case, alternatives of a decision) are “good” or “bad”, according to a given perspective (e.g. telling the truth is “good” according to ue of honesty). These evaluations can be compared with each other. Thus, things can be compared from the perspective provided by values (e.g. telling the truth is better than telling nothing from the perspective of honesty).

The second part is the value-system itself, combining values and indicating their relative importance on global evaluations. When evaluating a thing, evaluations of all values are combined together. This combination indicates the overall “goodness” of a thing (e.g. reporting a fault from my boss is “good” according to honesty and “bad” according to obedience). The centralizing point is the possibility to compare the global evaluations of two things. For instance, the decision “do not report a fault of my boss” is “neutral” according to honesty and obedience. This global evaluation can be compared with the one of “report the fault of my boss”. Since in my value-system “obedience” is more important than “honesty”, not reporting is “better” according to my value-system.

These models of value-systems can be used for making value-compliant decisions.

When making a decision, agents can raise multiple alternatives that can hardly be rationally solved (no obvious decision). In this case, agents can select the alternative that is the best according to their values.

We proposed an agent decision architecture that is both culturally-sensitive and capable of making decisions in situation of coordination. This architecture is based on the BDI architecture. This architecture is a traditional solution for performing practical reasoning, which is adequate for coordinating. This architecture is expanded for integrating the influence of value-systems. Throughout this chapter, we introduce in details possible influences of value-systems on internal decisions that are made within this architecture. Furthermore, we introduce and relate this architecture to four decision aspects which are both culturally-sensitive and influence coordination: identity/relationship, norms, uncertainty, and learning.

Towards Answering Research Question 3: How Can We Use Models of Human Culture as a Practical Tool to Support Coordination in Artificial Societies?

This chapter, in Section 4.1.2 raises a warning to take into consideration when using the influence of culture on values for supporting coordination in Chapter 7. Values supported by culture in human societies should not be blindly used assuming that they will necessarily improve coordination. Values supported by cultures generally but do not necessarily support coordination. Some of them can actually be detrimental for coordination (e.g. glorified self-sacrifice, proselytism). Thus system designers should be particularly careful about the range of values they integrate within their system.



5

Showing Evidence That Our Decision Architecture Can Replicate Properties of Human-Like Culture

Chapter Summary

This chapter aims at getting more confidence that (1) values are a core aspect of the influence of culture on coordination, and (2) our value-sensitive decision architecture is adequate for replicating human-like influences of culture on coordination.

In order to do so, this chapter highlights evidence that these two properties are fulfilled. More precisely, this chapter implements our decision architecture into an agent-based simulation where agents are sensitive to values. The outcomes of this simulation are tested against the core properties of the influence of culture on coordination as depicted by theories of culture. Furthermore, in order to get better confidence, we also show evidence that the match between our architecture and reality is credible with additional models based on our architecture.

In order to achieve the goal of this thesis, we need to recognize the *core*¹ aspects of the influence of culture on coordination and how to model this influence. These core aspects were investigated and modeled in Chapter 4. There, we showed that values are at the core of the influence of culture on coordination and we introduced a decision architecture for building value-sensitive decision models that can replicate the influence of culture in situations of coordination.

This architecture alone is insufficient with regard to our purpose: we need to be more confident that our architecture can actually replicate *in practice* the core manifestations of the influence of culture on coordination. So far, we only showed the compliance of this architecture with the mechanisms of the influence of culture on coordination as depicted by theories. This compliance from a theoretical standpoint gives some but only limited confidence that this architecture can actually replicate cultural manifestations in practice. Given the goal set by this thesis, we need to be more confident about the practical application of this model. Indeed, we aim at determining whether and how a model of the

¹This focus on these core aspects was further justified in the introductions of Chapter 3 and Chapter 4.

influence of culture can be used in practice for supporting coordination: we need to show that such a model exists.

In the current chapter, we aim at getting more confidence that the architecture presented in Chapter 4 can be used for building models that replicate in practice known manifestations of the influence of culture on coordination. What can we do for getting more confidence? Given that our architecture aims at replicating a natural phenomenon, we cannot “prove” that our architecture is correct. Nevertheless, we can increase our confidence by showing the relative correctness of our architecture, by gathering *evidence* for indicating that our architecture can be indeed used for building models that replicate core features of the influence of culture on coordination. The most direct evidence we provide in this chapter consists in using our architecture for building such a model and checking that the output of this model matches known influences of culture on coordination from human societies. More specifically, we build a simulation model of culturally-sensitive agents working in an IT department. Additionally, we dedicate a section for introducing further evidence (e.g. coherent results in spite of different culture, additional simulations).

Regarding other chapters, the current chapter relies on the theoretical background from Chapter 2 and the conceptualization from Chapter 3 for establishing a simulation scenario where the influence of culture on coordination can be expressed. Furthermore, this background and this conceptualization provide a set of expectations (e.g. the influence of culture on individual behaviors and collective performance) that are used for validating the output of our model. Following chapters rely on the validation presented in the current chapter by assuming that our model is relatively accurate, for elaborating social science theories in Chapter 6 and for exploring how to use such a model of culture can be used for supporting coordination in Chapter 7.

This chapter is structured as follows. We first present an informal overview for introducing the big picture of our simulation in Section 5.1. Then, we formalize concepts and static objects that structure our simulation model in Section 5.2. Next, in Section 5.3, we formalize the agent decision processes, based on our decision architecture from Chapter 4. Afterwards, we check whether our simulation model replicates the core properties of culture as depicted in theories of culture in Section 5.4. Finally, we propose additional evidence for supporting the accuracy of our architecture in Section 5.5.

5.1 Overview

This section introduces an overview of the concrete simulation we implement in the following sections. This overview aims at providing a general understanding of the core aspects of our simulation without being overwhelmed by implementation details. These details are further presented in the following sections.

As a statement of our intent, we aim at introducing a simulation (1) that introduces a human-like multi-agent coordination setting and (2) where culture can be introduced and meaningfully influence decisions of agents. In particular with (1), we want to be able of measuring the basic information about individual behaviors and collective outcomes (e.g. performance depending on the environment). By using (1) and (2), we can check whether the combination of culture and coordination outcomes match expectations provided by theories of culture.

This overview is introduced in two parts. Section 5.1.1 introduces the environment and the coordination setting of our simulation (e.g. coordination mechanisms in place, decisions that agents have to make). Section 5.1.2 introduces more specifically core aspects of agent decision processes.

5.1.1 Environmental and Coordination Background

Our simulation replicates the work of a Department Of Information Technology (DO-IT). DO-IT is an organization that aims at providing technical support to its clients.

DO-IT provides three kind of services: *maintenance* (e.g. keeping email servers running), *user support* (e.g. helpdesk) and *software development* (e.g. providing code for other organizations). From an organizational standpoint, the *maintenance service is more important than the other ones*. Failing to provide this service leads to the strongest user dissatisfaction, impacting in turn to organizational performance (no bonus, less trust).

For getting access to these services, DO-IT receives requests from its clients. Clients contact the organization, indicating a support they need. If the support is provided in due time, the task has *succeeded*. Otherwise, when the client runs out of patience, he or she calls another service and the task has *failed*, impacting on organizational performance.

Internally, requests are treated as independent *tasks*. These tasks are shared in a common workspace. Each task is open as a contract, which can be accepted by any employee and thus removed from the shared workspace. We consider two possible *organizational structures* in order to investigate various properties depicted by theories of culture². DO-IT can either be an *adhocracy* or a *machine bureaucracy*. In adhocracies, employees are free to select which contract they accept. In machine bureaucracies, a *leader* is in charge of coordinating agents. Leaders “strongly suggest” contracts to subordinates depending on their expertise. Then, subordinates remain free to select any available task, disobeying or not.

Once a task is allocated, agents have to work on them until completion. The duration for solving a task varies depending on multiple parameters. It depends on *how* this task is solved: *quickly, creatively* (providing an original ad-hoc solution) or in relying on *standards*. The duration also depends on whether the *expertise* of the solving agent matches required expertise for tackling the task.

5.1.2 Decision Aspects

This setting forces agents to make decisions that impact coordination. When they are available, they have to decide which contract they accept and how to solve this contract (quickly, creatively or in relying on procedures). Multiple aspects influence these decisions.

Some aspects depend on the environment and the coordination structure. In particular, the type of available tasks and the match between the expertise of the agent and the expertise required for handling the task.

²We selected the two formal organizational structures that match the two cultures we investigate in this chapter. More details about this match is presented in Figure 2.3. Furthermore, these two organizational structures match very different profiles (adhocracies handle better complex and dynamic environments while machine bureaucracies handled better simple and static environments), as introduced in Section 2.1.2.

Some other aspects are internal to the agent. Tasks can satisfy more or less the *motives* of the agent (survival needs, need for conformity and need for self-expression). Decisions can be related to aspects introduced in Section 4.4 (culturally-sensitive aspects). In particular, agents have to take into consideration norms (e.g. being conform to others), uncertainty (e.g. managing tasks with unknown deadlines) and identity/membership (e.g. complying with orders of leaders).

These internal aspects are connected to a rich conceptual background that needs to be further introduced. Furthermore, culture can impact the influence of these aspects on decisions. This section is dedicated to further introducing further the influence of these aspects on decisions.

Culture

In this model, we consider the influence of two values: *conservatism* and *openness to change*, drawn from Schwartz (1999)³. These values are related in Section 3.2 to two HCDs: PDI-CD and UAI-CD. Conservatism is conceptually related to high PDI-CD and high UAI-CD and openness-to-change to low PDI-CD and low UAI-CD.

In this model, we kept the influence of values as simple as possible. They only influence the relative importance given to the various motives. More precisely, the value of conservatism supports the motive of conformity and the value of openness-to-change supports the motive of self-expression. Nevertheless, this influence is sufficient for influencing behaviors of agents coherently with regard to expectations from social sciences (e.g. conservative agents tend to perform procedural resolutions while agents which are more open-to-change tend to perform more creative resolutions).

Motives

Motives constitute one of the core cognitive mechanisms that push individuals (and their agents models) to action. So far, motives were not extensively introduced in this thesis. They are relatively human-like decision aspects that have not been extensively used for handling coordination problems. Nevertheless, this human-like aspect actually makes them particularly useful in this chapter.

Motives are integrated within this simulation because they are both useful for supporting human-like decisions in situation of coordination as illustrated by Norman and Long

³Why these values?

We preferred to focus on a low number of values. Indeed, values are tightly embedded within decisions. Thus adding more values is likely to increase model complexity (this is what happened when we added the values of achievement and benevolence). Instead, we aim at keeping a simulation model as simple as possible.

Why these two values? In order to select these values, we rely on the main nine clusters of values, recognized by Schwartz (1999), presented in Figure 4.1, looking for abstract, central values. We select these two values because they particularly relate to the extremes of PDI-CD and UAI-CD. These extremes, in turn, relate to different and opposed Mintzberg organizational canonical structures. These structures have in particular opposed ways of conceiving interactions (e.g. let leaders decide vs. let subordinates decide; rely on standards vs. rely informal creative resolutions). Thus, these values provide a solid ground for raising sharply different forms of culturally-driven preferred coordination patterns. In turn, the preferred coordination patterns match to (or oppose) the organizational structures we study in this chapter (machine bureaucracy or ad-hocracy), as depicted in Figure 2.3.

All is set for raising visible oppositions, as depicted by theories of culture!

(1996) and are sensitive to culture. In addition, motives offer an interesting parallel with values which initially drove us towards exploring this aspect.

Basic Properties of Motives: Motives indicate basic desires or needs of individuals, pushing them to action (e.g. hunger, thirst, need for love, need for recognition). Needs can be more or less satisfied. The least a need is satisfied, the more individuals are driven towards fulfilling it. The satisfaction of each need evolves depending on the situation and the actions of individuals.

As a core property, motives can be partially ordered in terms of importance. Fulfilling unsatisfied more important motives drives more attention on decisions than fulfilling unsatisfied less important motives (e.g. breathing is more important than eating which is more important than social recognition; considering this, social recognition is generally discarded when an individual cannot breathe). When satisfied, the influence of a motive on decision fades out. Thus, less important motives are influential when more important motives are satisfied (e.g. if breathing is secured, individuals can start acting towards getting food). In this direction, Maslow argued that there exists an inherent order within motives, as illustrated in Figure 5.1. In particular, survival-related motives are given more importance than others.

Motives and Culture: Hofstede et al. (2010a) explore the influence of culture on motives. Lowest-level motives tend to be more related to human nature and thus independent on culture. Conversely, higher-level motives tend to be more sensitive to culture. Hofstede et al. (2010a) explain that culture alters the order of higher-level motives (e.g. need for esteem can become less important than need for belonging).

Motives offer an interesting point of comparison with values. Motives are “motivational”, they can be seen as desires. They push individuals towards making decisions⁴, either directly towards taking action or indirectly towards adopting goals. Conversely, values are beliefs telling what is “good”. Given a decision to make, they indicate which alternatives are preferable. In some sense, motives push to make decisions and values help making these decisions.

Motives in Our Simulation: In our simulation, motives are used for modeling the desire update function, as presented in Chapter 4. They indicate what agents want and how much they want it.

In this chapter, we consider three motives: survival motive, conformity motive and self-expression motive⁵. Survival motive is a low-level motive that is thus prioritized by agents. This motive is satisfied when the agent believes that its job is stable. More concretely, this motive is satisfied when the organization performs well. Conformity motive is satisfied when agents solve tasks according to standards and otherwise becomes unsatisfied with time. Self-expression motive is satisfied when agents solve tasks creatively and otherwise becomes unsatisfied with time.

Culture, expressed through the value of conservatism and openness-to-change, influ-

⁴Motives are not the only reason for making decisions as individuals can also react to the environment. Nevertheless, motives push towards pro-activeness, raising decisions

⁵We select these three motives for two reasons. First, we want to differentiate explicitly the influence of culture from the influence of human nature (survival motive vs. others). Second, we want to have motives influenced by conservatism (conformity) and openness-to-change (self-expression).

We rely on these three motives and these direct relationships with culture for sake of simplicity. This model can be expanded with more motives and with more complex links (e.g. tying conformism with survival motive and conservatism motive).

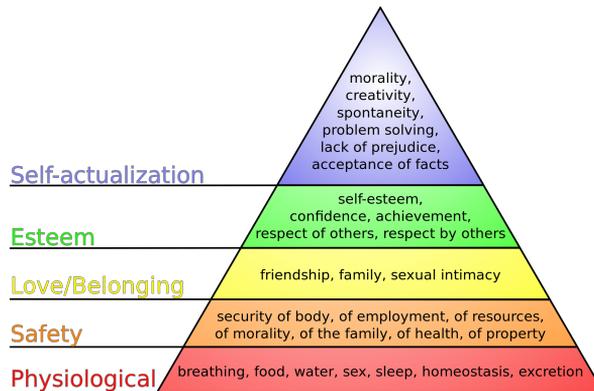


Figure 5.1 – Maslow Pyramid of Needs (Maslow (1943))

ences the relative preference between high-level values (conformity and self-expression). The more agents prefer conservatism, the more the conformity motive will matter for them. Conversely, the more agents prefer openness-to-change, the more they will favor the satisfaction of self-expression. Culture does not influence survival, which is more related to human nature.

Uncertainty & Norms

The setting of this simulation raises the two aspects of uncertainty and norms, as introduced in Section 4.4.2 and Section 4.4.3. In order to handle uncertainty, agents can use standards for solving tasks. Furthermore, agents may want to behave like other agents (social norms). Culture influences through motives the relative importance of these two aspects. Depending on the (cultural) importance given to conservatism, agents are more or less prone to follow protocols and to behave as others.

Identity/Membership

The setting of this simulation raises decisions where identity is important. In particular, the machine bureaucracy features the roles of leaders and subordinates. Depending on culture, agents can be more or less obedient with regard to orders provided by the leader.

Conclusion

The overview presented in this section matches the requirements for testing our model. We introduced a framework for representing a variety of cultural expressions, which can have a broad range of influences on multiple aspects of decisions made by individuals. Our simulation also allows measuring collective outcomes (e.g. organizational success, overall satisfaction).

In conclusion, this overview depicts a range of models that can be used for testing our architecture. Thus, we can check whether the output of these models matches known influences of culture on individual decisions and collective outcomes, as observed in reality

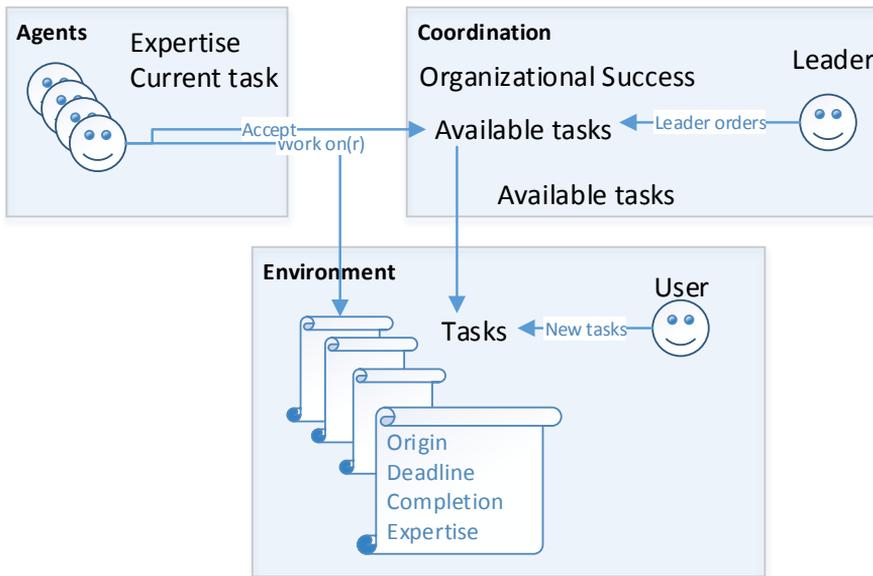


Figure 5.2 – Graph introducing the core concepts used by this simulation and the various entities (agents, objects). Links indicate possible interactions between entities

5.2 Simulation Model

This section formalizes the general setup introduced in the overview (Section 5.1). This formalization introduces the specific concepts, entities (objects and agents), influences in terms of coordination and the dynamics of this simulation. These concepts and entities depict the environment in which agents will have to make decisions.

The structure of this section is as follows. Section 5.2.1 introduces our model of the environment (entities, objects, agents). Section 5.2.2 introduces the coordination model for handling this environment. Finally, Section 5.2.3 introduces the update loop that formalizes how the various entities of the environment evolve with time.

5.2.1 Environmental Model

Environment in MASs are composed of entities (e.g. objects, bodies of agents). These entities are static. They remain within the environment throughout the simulation (unless being explicitly removed). They can evolve with time or as a consequence of actions of agents, as formalized by the simulation loop in Section 5.2.3. The current section introduces entities populating the environment and their properties, based on the general

overview presented in Section 5.1.

Our environment is composed of two kinds of entities: tasks and agents. Their properties are introduced in the following sections. In addition, the environment keeps track of the current time. Formally the time is represented by the variable $time \in \mathbb{N}$. Each execution of the update loop represents the occurrence of one time unit, modeling one hour of working-time.

Tasks

Tasks represent the requests sent by clients to DO-IT. The set of tasks is represented by T . Tasks have several basic properties.

Deadline: clients expect DO-IT to solve tasks in due time. Otherwise, clients run out of patience, withdraw their demand and complain about DO-IT to business-ranking websites. The patience of the client is represented by a deadline, formalized by $deadline: T \rightarrow \mathbb{N}$. For any task $t \in T$, $deadline(t)$ is the time at which the client cancels the request and complains. In the update loop, this deadline is compared against $time$ for determining whether a task has failed or still ongoing. Task deadlines are not visible: this information is private to the client. Agents cannot access this variable for considering how much time they have remaining.

In our simulations, deadlines occur 20 rounds after tasks are issued. Formally, if a task $t \in T$ is issued at time $t' \in \mathbb{N}$ then $deadline(t) = t' + 20$. This duration gives enough time for experts to process easily these tasks while introducing risks when attempting inappropriate resolutions, particularly if the solving agent is not an expert.

Origin: DO-IT offers three types of services: maintenance, support and software development. Maintenance tasks are the core activity of the organization and their achievement is the most important with regard to organizational success (introduced in Section 5.2.2). Support and development tasks have similar features in the current model and are equivalent here. Nevertheless, they were used for expanding the content of the model, as introduced in Section 5.5.

Formally, the set of possible origins of tasks is represented by $O = \{o_m, o_s, o_d\}$, where o_m stands for maintenance tasks origin, o_s for support tasks and o_d for software development tasks. The origin of a given task is represented by $origin: T \rightarrow O$. In our simulations, the origin of a task is randomly drawn from O .

Expertise: Tasks can be completed more or less efficiently depending on the match between agent's expertise and the expertise required for solving this task. Formally, the set of possible expertise is represented by E . The expertise required for solving a task more efficiently is represented by $expertise: T \rightarrow E$. In our simulations, the expertise required by a task is uniformly randomly drawn from E . Thus, since agents have a bounded expertise, $|E|$ indicates the relative difficulty for having the adequate expertise for handling incoming tasks. In other words, $|E|$ indicates the relative complexity of the environment.

Progress: Tasks are not performed in a single time-unit. To that extent, we need to record the amount of progress that remains to be performed on a given task. Formally, the degree of progression of a task is represented by $completion: T \rightarrow [0, 1]$. A completion of 0 means that the task has not been started at all and a completion of 1 means that the task has been successfully completed. This progression evolves when agents work on it, as further introduced in the update loop.

Properties of Agents

Agents have some properties that are introduced now for keeping the model simple. From a purely theoretical perspective, these properties are part of agent decision process, but making them external to the agent allows simplifying a lot of representations and decisions of agents.

Expertise: Agents can be experts in a domain, influencing how efficient they are for solving some tasks. Formally, this expertise is represented by $expertise_a \in E$. If an agent a is expert in $e \in E$ if $e = expertise_a$. An agent a works on a task t such that $expertise(t) = expertise_a$, then t will be completed faster than if $expertise(t) \neq expertise_a$. More details about this relation are introduced in Section 5.2.3. In our simulations, $expertise_a$ is randomly drawn from E .

Current Task: Agents are allocated a task when they decide to accept a task. Then, agents work on it until it is ended (succeeded or failed), as presented in the update loop in Section 5.2.3. Formally, the task that is allocated to an agent a is represented by $allocated_a \in (\perp \cup T)$, where \perp means “no task”. If $allocated_a \neq \perp$ then a is working.

solving Tasks: When solving a task, agents adopt a general strategy⁶ for determining how to tackle it. We consider three strategies: quick fixing, relying on procedures (or standards) and creative/adaptive action (e.g. using intuition or locally appropriate solutions). Formally, the set of task resolutions is $R = \{r_{qf}, r_{stand}, r_{crea}\}$. Quick fixes are represented by r_{qf} , resolutions following procedures are represented by r_{stand} and creative resolutions are represented by r_{crea} . The strategy selected by a is represented by $resolution_a \in R$. This strategy influences the performance of the resolution as indicated in Section 5.2.3 but also the satisfaction of the agent as introduced in Section 5.3.

5.2.2 Coordination Structure

DO-IT is modeled by an organization. This organization can be characterized by three main features: the allocation system, its structure and the organizational success.

Allocation System: DO-IT features a computer system that records the set of tasks to be solved and can allocate them to agents. The set of tasks to be solved is the set of tasks that have already been proposed by users, which are not failed and not currently allocated to an agent.

From a modeling perspective, this set of tasks is represented by $available \subseteq T$ which is visible by all agents. When an agent asks for a task, this task is removed from this set and allocated to the agent, as presented in Section 5.2.3.

Organizational Structure: DO-IT can have two types of structures, defined at the setup of the organization.

The first structure is an adhocracy. In this structure, agents do not have formal relations with each other. The second structure is a machine bureaucracy. In this structure, a leader agent indicates to subordinates which task they should perform. Subordinates may accept it or disobey and perform another task. This message is formally represented by

⁶We acknowledge that, theoretically speaking, these strategies should be part of the agent decision process (e.g. plan). Nevertheless, by making it external, we can pinpoint more easily to when agents make the decision of their general strategy. Plus, we do not need to go into within the details of agent decisions for explaining the update cycle.

Algorithm 1 Simulation Update Cycle

```

Initialize agents, the organization and the environment
for the desired number of iterations do
  Increment time
  Add new requests from clients to available
  for all agt ∈ A which is not working do
    if Any available task? then
      Allocate decision(select – task(), agt) to agt
      Set the resolution strategy to decision(select – strategy(), agt)
    end if
  end for
  for all agt ∈ A which is working do
    Work on tasks
  end for
  Manage ended tasks
end for

```

$t_l \in available$. For sake of simplicity, we only refer to this agent by using t_l . Thus, we avoid to introduce an extensive decision model for the leader.

Organizational Success: DO-IT can be more or less successful. This success depends on the amount of tasks successfully performed lately. In particular, maintenance tasks, the core mission of DO-IT, have more impact than other tasks on the organizational performance.

From a modeling perspective, organizational success is represented by $OS \in [0, 1]$. Its dynamics are presented in Section 5.2.3.

5.2.3 Update Loop

The update loop describes how the entities within the simulation evolve with time and when agents have to make decisions. This loop is described by Algorithm 1. Each cycle of the update loop simulates one hour of working-time.

The rest of this section further introduces what happens in each step.

Initializing Agents, the Organization and the Environment

The initialization phases generates agents, sets up the coordination model and the environment. *time* is set to 0.

The organization is setup as follows. The organizational structure is determined by the simulation setup, enabling the use of t_l or not. The range of available tasks *available* is set to \emptyset . The organizational success *OS* is set to 1.

Fifteen agents are created. Each agent is randomly allocated one expertise from *E* and its current task is \perp . The number of agents is determined for scaling with the amount of incoming tasks and the time required for processing them. In our experience, adding more agents did not lead to significant changes in terms of outcomes.

Adding New Requests from Clients

DO-IT receives new calls from clients. Each iteration, DO-IT receives five new tasks. These tasks are added to the set of available tasks *available*. We set this number so the organization is not overloaded by the number of incoming tasks but high enough in order to keep providing work for agents and forcing agents to select adequate tasks and resolutions.

The origin and the expertise of incoming tasks is determined randomly, the deadline of the task is set to $time + 20$ and its *completion* is set to 0.

Decide on Which Task to Accept

Agents that are not currently working (for which $allocated_a = \perp$) consider the range of available tasks (if any) and decide on which task they want to be allocated. The set of available actions for this decision is $\{accept(t) | t \in available\}$.

When making this decision, agents receive a set of percepts. Agents perceive the set of available tasks *available* and the organizational success *OS*. If the organizational structure is a machine bureaucracy, each agent receives a message from the leader suggesting which task to work on, through $t_l \in available$.

The task t_l that the leader suggests to subordinates is determined as such. Leaders try to allocate the subordinates tasks that match the best with their field of expertise. Otherwise, they allocate them a random task. Formally, the proposed task for an agent a is a random task from $\{t \in available | expertise(t) = expertise(a)\}$ if this set is not empty. Otherwise, the task is randomly drawn from *available*.

Once an agent a selected an action of the form $select(t)$, then t is removed from *available* and $allocated_a \leftarrow t$.

Decide on a Resolution Strategy

When receiving a new task, agents have to decide on which resolution strategy they rely on for processing this task. The range of available actions for this task is $set_resolution_strategy(r), \forall r \in R$. If an agent a decides on the action of the form $work(r)$, then $resolution_a$ is set to r .

Working on Tasks

Each round, each agent a that has an allocated task t and a resolution strategy r works and progresses on t . Formally, this progress is represented by an increase of $completion(t)$ by $progress(t, a, r)$:

$$progress(t, a, r) = \frac{expertise_match(t, a) \times resolution_influence(t, a, r)}{task_duration(t)}$$

where $expertise_match(t, a)$ represents the influence of expertise on agent efficiency, $resolution_influence(t, a, r)$ represents the influence of the selected resolution on agent efficiency and $task_duration(t)$ is the base duration of task t .

$expertise_match(t, a)$ represents the increase in individual performance when solving the adequate task with the required expertise. $expertise_match(t, a)$ is represented by:

$$expertise_match(t, a) = \begin{cases} 1 & \text{if } expertise(t) = expertise_a \\ 0.5 & \text{otherwise} \end{cases}$$

$resolution_influence(t, a, r)$ represents the influence of the selected resolution on progress. Quick-fix resolutions are the fastest. Procedural resolutions can be fast if a possesses the adequate expertise, otherwise these resolutions are much slower. Finally, creative resolutions are more adaptive, being slower than quick-fixes, but faster than the application of a procedure when lacking the adequate expertise. Formally, $resolution_influence(t, a, r)$ is represented by:

$$resolution_influence(t, a, r) = \begin{cases} 1 & \text{if } r = r_q \vee (r = r_{stand} \wedge expertise(t) = expertise_a) \\ 2/3 & \text{if } r = r_{crea} \\ 1/4 & \text{otherwise } (r = r_{stand} \wedge expertise(t) \neq expertise_a) \end{cases}$$

Finally, $task_duration(t)$ represents the base task resolution, which is set to 3 in our simulations.

These values are balanced such that the processing power of the organization matches with incoming tasks and task duration. Furthermore, we represent a performance increase when agents solve tasks they are expert at. Finally, inappropriate resolutions take more effort, triggering failures.

Manage Ended Tasks

Once all agents have worked on their tasks, we check whether tasks are completed or failed. We consider all tasks in a random order. For each task t if $completion(t) = 1$, then t solved⁷. In this case, OS is updated as follows:

$$OS \leftarrow OS + (1 - OS) \times \begin{cases} 0.02 & \text{if } origin(t) = o_m \\ 0.01 & \text{otherwise} \end{cases}$$

Otherwise, if $time = deadline(t)$ then the task has failed. In this task OS is updated as follows:

$$OS \leftarrow OS \times \begin{cases} 0.95 & \text{if } origin(t) = o_m \\ 0.99 & \text{otherwise} \end{cases}$$

If t has failed or succeeded, then t is removed from *available* if t is in this set. If t is allocated to an agent a , then $allocation_a = \perp$ and a is notified that t is finished.

Conclusion

This section introduces the entities populating our simulation and their dynamics. Furthermore, this section presents the decisions that agents will have to make: selecting tasks and the strategy for processing tasks.

⁷For some resolutions, $completion(t)$ can exceed 1. In that case $completion(t)$ is set to 1

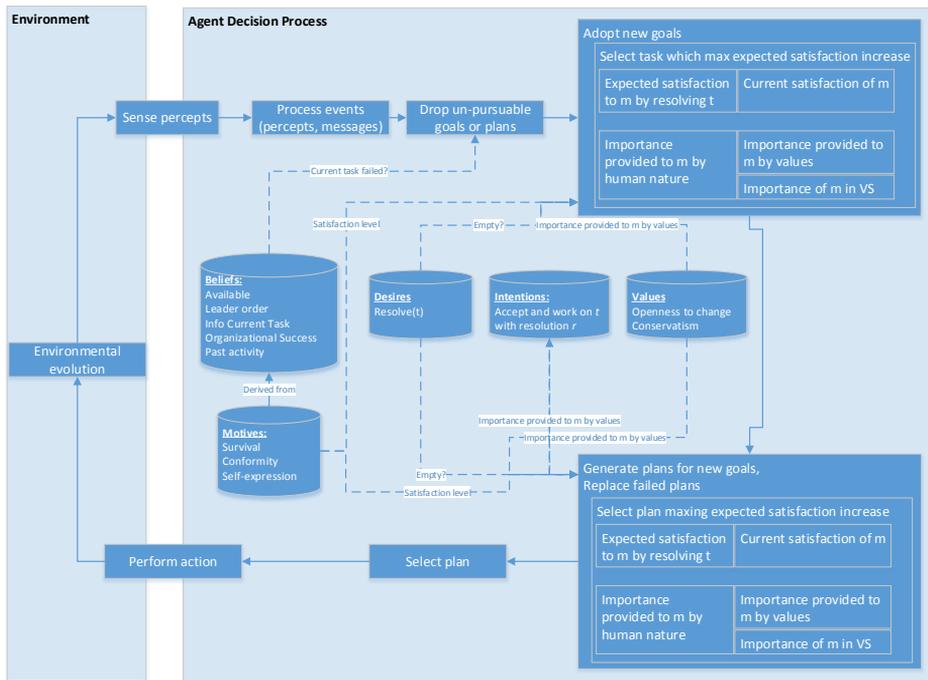


Figure 5.3 – Core elements of agent decision process. Plain lines represent the decision loop. Dotted lines represent interactions with agent memory.

This coordination and environmental model offers numerous opportunities for expressing cultural values that we explore in this chapter (namely, conformity and openness-to-change). Will agents prefer to solve their tasks according to good old standards or more creatively? What about when the organization will start to fail, will they keep doing what they (culturally) prefer or go for maintenance tasks? What if there is a leader watching them?

5.3 Agent Decision Process

As presented in Section 5.2.1, agents have to make decisions: determining which tasks they want to work on and determining the general strategy for achieving their tasks.

Our agent decision process instantiates the decision architecture presented in Section 4.3. This architecture is composed of a memory structure presented in Section 5.3.1 and an internal decision loop, presented in Section 5.3.2. An overall depiction of the decision process is proposed in Figure 5.3

Within these standard decision aspects, we express an additional specific decision aspect: motives or needs. This aspect represents growing needs of individuals, influencing decisions in multiple points as presented in Section 5.1. *These motives are also the en-*

try point for integrating the influence of culture on decisions. In this simulation, we consider three motives that are used throughout this section: survival, conformity and self-expression. These needs are modeled by $M = \{m_s, m_c, m_{se}\}$ where m_s is the motive for survival, m_c for conformity and m_{se} for self-expression.

5.3.1 Memory Structure

The memory structure represents the static elements that belong to the agent. As described in Chapter 4, this structure is composed of four blocks: beliefs, desires, intentions and value-system.

Beliefs

Beliefs turn percepts into (durable) information about the current situation.

Environmental Beliefs: Beliefs record information about the state of environment and about other agents: *available* the set of available tasks, t_l the order given by the leader (if any), *allocated* indicating which task is allocated to the agent (if any), *resolution* indicating the resolution strategy selected by the agent and *OS* indicating the organizational success. From a modeling perspective, since these beliefs are accurate, we use the same symbols as those used for depicting environmental properties.

Beliefs for Representing Motivation: In addition to these externally-given beliefs, agents have special belief structures for recording and measuring their satisfaction. Formally, the satisfaction of a motive of an agent a is represented by the function $sat_a: M \rightarrow [0, 1]$ where $sat_a(m)$ measures the satisfaction of motive $m \in M$ of agent a .

The *survival motive* m_s is, in this simulation, related to the expectation that the job of the agent is safe. This drive is indirectly related to the organizational success: if DO-IT performs badly it may close, ending the job of the agent. Formally, the satisfaction of the survival drive m_s is represented by:

$$sat_a(m_s) = \frac{1}{1 + e^{-4(2OS-1)}}$$

This formula links m_s with OS by using a sigmoid function, represented in Figure 5.4. This sigmoid function aims at representing that, when satisfied, survival drives are easily blurred. However, if this drive is no longer sufficiently satisfied (when OS is below 0.7), then this drive becomes quickly a great source of concerns.

The satisfaction of the *conformity motive* is related to the recent compliance with regard to standards. Formally, the satisfaction of this motive is related to belief $sat_a(m_c) = past_activity_c$. This later variable records whether the agent behaved in conformity with standards, discounted with time. More information about the update of this variable is introduced in Section 5.3.2.

The satisfaction of the *self-expression motive* is similar to the conformity motive, except that it depends on whether the agent has recently performed creative resolutions. Formally, the satisfaction of this motive is represented by the variable $sat_a(m_{se}) = past_activity_{se}$, which records whether the agent recently behaved creatively.

Desires

Desires record the concrete goals that are currently pursued. In this model, the set of goals that an agent can adopt is $solve(t)$ for any task $t \in T$. The determination of such a goal occurs when the agent has to decide which task to select, in the decision $select - task()$. This decision is strongly influenced by motives as presented in Section 5.3.2.

Intentions

The set of plans which can be adopted as intention has the form: “accept a task; select a resolution; perform the task until the task is ended”. The determination of such an intention occurs when the agent decides which task to select, in the decision $select - task()$. The decision $select - strategy()$ is thus simplified, since it only executes the second step of this plan. The selection of a given intention supports the type of actions that are to be performed, which is crucial for satisfying motives (e.g. creative resolutions).

Value System

Chapter 4 proposes a decision architecture for modeling the influence of values on decisions. This architecture relies on a two-step process: first defining values and then combining them together.

Values: In this simulation, we consider two values introduced by Schwartz (2006b): conservatism/tradition and openness-to-change. These two values are represented by $V = \{v_c, v_{otc}\}$. Values evaluate the “goodness” of a decision from a limited perspective.

Chapter 4 introduces a more specific framework for modeling values. In this framework, each value $v \in V$ is modeled by using three items. (1) the set of possible evaluations $eval_v$, (2) a partial order over this set $<_v$ for indicating which evaluations are better than others, (3) for each decision d to be made, an evaluation function $evaluate_{v,d} : Alt_d \rightarrow eval_v$ which indicates how “good” are the various alternatives for a decision d (the set of alternatives is Alt_d) from the perspective of value v .

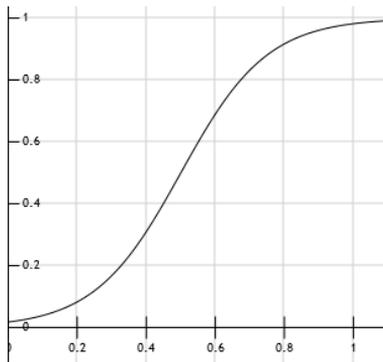


Figure 5.4 – Plot of $sat_a(m_s) = \frac{1}{1+e^{-4(2OS-1)}}$, the sigmoid function used for evaluating the survival motive, based on organizational success OS

For both values, we represent $eval_v$ with the range $[0, 1]$ and $<_v$ with the standard order for real numbers. With this representation 0 means “very bad” and 1 means “very good” from the perspective of this value. This representation matches well with the numerical representations that are proposed throughout this simulation (e.g. satisfaction of motives are also modelled with real numbers).

Evaluation functions are sensitive to concrete decisions to be made. Thus, they are described later, when concrete value-sensitive decisions have to be made.

Value-Systems: Values are combined together for defining a value-system capable of determining which alternative of the decision is the best supported by values.

In this direction, Chapter 4 introduces a solution for designing such a combination, in ordering tuples of the form $eval_{v_1} \times \dots \times eval_{v_n}$ with the $<$ operator. Then, given a set of alternatives Alt_d , the preferred alternative is one of the maximal elements of Alt_d according to $<$.

In this simulation, we represent $<$ by using the following weighted sum:

$$evaluate_d(alt) = \alpha \times evaluate_{v_c,d}(alt) + (1 - \alpha) \times evaluate_{v_{otc},d}(alt)$$

With this sum, α weights the relative importance between openness-to-change and conservatism. The higher is α , the more importance is given to conservatism.

As an illustration, consider the decision of writing a piece of a Java program. After a bit of reasoning, the agent faces two rationally relatively-similar alternatives that values can help separating. (1) rely on a few good old design patterns or (2) create an object that is specifically attuned for this function. Solution (1) is a traditional yet extremely boring solution for handling this problem. This alternative is in line with tradition $v_c = 0.7^8$ but is not at all open to change $v_{otc} = 0.1$. Solution (2) is neither the standard supported by the organization $v_c = 0.5$ neither particularly creative $v_{otc} = 0.5$. In a very conservative culture, $\alpha = 0.8$, the first alternative is evaluated as 0.58 while the second alternative is evaluated at 0.5. In this culture, better rely on the good old yet boring solution (1). In a different culture, which is more open to change (e.g. $\alpha = 0.5$), solution (2) is preferred.

This representation has numerous advantages. It is relatively comprehensive and easy to control with α . This factor also offers the possibility of making softly evolving α which is helpful for exploring in details the influence of various value-systems on decisions when running simulations. Furthermore, this representation matches well the soft representation of the evaluation function we proposed for each value.

5.3.2 Decision Loop

The decision loop determines a sequence of internal decisions to perform in order to determine which external action(s) to perform and update the internal memory of the agent. The decision loop introduced in this section is a direct implementation of the decision loop presented in our architecture (see Figure 4.2).

Process Percepts

As described in Section 5.2.3, agents perceive accurate information about *available*, *OS*, *t₁*. They are also warned through an event when the task were working on just ended (succeeds

⁸We use a reader-friendly representation. This notation refers to the evaluation function

or failure). All these pieces of information are recorded by agents. In addition, agents update two variables that record their recent activity: $past_activity_c$ and $past_activity_{se}$, according to the following formula. The evaluation the satisfaction of agents' motives is later derived from these variables

$past_activity_c$ is updated according to this formula, where r is the current resolution selected by the agent:

$$past_activity_c \leftarrow \begin{cases} past_activity_c + 0.05 \times (1 - past_activity_c) & \text{if } r = r_{stand} \\ past_activity_c \times 0.99 & \text{otherwise} \end{cases}$$

$past_activity_{se}$ is updated according to this formula, where r is the current resolution selected by the agent:

$$past_activity_{se} \leftarrow \begin{cases} past_activity_{se} + 0.05 \times (1 - past_activity_{se}) & \text{if } r = r_{crea} \\ past_activity_{se} \times 0.99 & \text{otherwise} \end{cases}$$

These two formula have two core properties we want to model: (1) satisfaction increases or decreases when the relevant resolution is used or not (2) the importance of past events is discounted over time.

Drop Goals or Plans

Agents drop the goal $solve(t)$ if they received an event indicating that t to be ended (finished or failed).

Adopt New Goals

If a goal is already currently pursued or no task is available, nothing happens.

Otherwise, agents adopt the goal $solve(t)$ where $t \in available$ is the available task for which the accomplishment is expected to maximize the satisfaction of the agent's motives. Formally, the formula for determining which goal (of the form $solve(t)$) is adopted is:

$$t = \operatorname{argmax}_{t' \in available} \operatorname{expected_satisfaction}(t')$$

where $\operatorname{expected_satisfaction} : T \rightarrow [0, 1]$ is the satisfaction expected to be provided by solving a given task. Since this satisfaction is dependent on the resolution strategy that will be selected by the agent, then $\operatorname{expected_satisfaction}$ can be expressed as:

$$\operatorname{expected_satisfaction}(t) = \max_{r \in R} \operatorname{expected_satisfaction}(t, r)$$

where $\operatorname{expected_satisfaction} : T \times R \rightarrow [0, 1]$ is the expected satisfaction provided by solving task t with resolution r .

As introduced in Section 5.1, the satisfaction expected to be provided by a decision depends on the various motives satisfied by this decision. Nevertheless, the various motives do not have an equal influence at any time. Sometimes, some motives have more impact on decisions than others. Three core aspects can be considered for conceptualizing this influence. First, the relative *need level*. The least a motive is satisfied the more this need

impacts on decisions. Second, the *expected satisfaction* that this decision will offer to this need. The more a decision satisfies a need, the more interesting this decision is. Third, the *relative importance* between needs. The satisfaction of some needs is always preferred over other needs all other things (need level and expected satisfaction) being equal. Formally, this influence is represented in our simulations by:

$$\begin{aligned} \text{expected_satisfaction}(t, r) = & \sum_{m \in M} \text{need_level}(m) \\ & \times \text{expected_satisfaction}(m, t, r) \\ & \times \text{importance}(m) \end{aligned}$$

where $\text{need_level}: M \rightarrow [0, 1]$ indicates the degree of need of a motive, $\text{expected_satisfaction}: M \times T \rightarrow [0, 1]$ is the expected satisfaction provided by performing task t with resolution r for motive m and $\text{importance}: M \rightarrow [0, 1]$ is the importance of motive m . These three aspects are introduced in turn in this section.

This formula balances the three core aspects of motives presented in Section 5.1.2 and leads to an influence that matches expectable properties from the influence of motives. More important motives are given higher priority in general. When they are sufficiently satisfied then less important motives start to impact decisions more and more. A motive which is not satisfied at all tends to drive more attention when making decisions and thus to be satisfied soon. Opportunities for satisfying multiple motives at once tend to be given more importance.

Need Level: The need level represents the degree of dissatisfaction of a motive. Given the normalization over $[0, 1]$ of all formula, we can formally represent the need level with:

$$\text{need_level}(m) = 1 - \text{sat}_a(m)$$

Importance of Motives: As introduced in Section 5.1.2, a core property of motives is that they are not all equally important. Some motives are more important than others: if two motives are equally dissatisfied and can be equally satisfied, more important motives are selected first. More generally, more important motives (e.g. survival-related motives), as long as they are not satisfied, tend to blur completely the influence of less important motives on decisions (e.g. being creative).

Another property to take into consideration in our model is the influence of culture. Culture only influences the importance of higher-level motives, while lower-level motives are more dependent on human nature.

In our model, the survival motive should be more independent from culture than the motives of tradition and openness to change. For sake of simplification, we assume that the survival motive is independent from culture while the other two motives depend only on culture.

Formally, we represent the importance of motives as such. $\text{importance}(m_s)$ is set to 0.7, $\text{importance}(m_c)$ is set to $0.3 \times \alpha$ and $\text{importance}(m_{otc})$ is set to $0.3 \times (1 - \alpha)$. Note that for keeping a general coherence within this model and avoiding introducing hidden complexity within formula, the sum of the importance of all values is normalized to 1.

With this formula, the survival motive is both culturally independent and given the highest importance. Then, culture influences the relative importance between conservatism and openness to change.

Expected Satisfaction: The expected satisfaction provided by performing a task depends on the resolution strategy that the agent will apply. Since the agent is assumed to maximize its satisfaction, $expected_satisfaction(m, t)$ is expressed as:

$$expected_satisfaction(m, t) = \max_{r \in R} expected_satisfaction(m, t, r)$$

where $expected_satisfaction(m, t, r)$ is the satisfaction expected to be provided by solving t with resolution r with regard to a motive m .

Multiple aspects are at play when considering the satisfaction provided by the resolution of a task with a given resolution.

First, general practical aspects. Inefficient resolutions are expected to be less satisfactory. Formally, the expected satisfaction is represented by:

$$expected_satisfaction(m, t, r) = progress(t, a, r) \times expected_satisfaction'(m, t, r)$$

where $expected_satisfaction'(m, t, r)$ is the base expected satisfaction provided by the completion of a plan and $progress(t, a, r)$ represents the relative efficiency of the agent for handling this task, represented by progress performed in one iteration when tackling the task, as introduced in Section 5.2.3.

The base expected satisfaction $expected_satisfaction'(m, t, r)$ depends on the nature of the task, whether this task is suggested by the leader (depending on the organizational structure) and the selected resolution strategy. Satisfaction of survival motive m_s depends on organizational structure. In machine bureaucracies, leaders tend to be empowered and be in charge of managing organizational success. Thus, in machine bureaucracies, compliance with orders is strongly related to the expected satisfaction of survival drives. In adhocracies, everyone is equally responsible for organizational good performance. In this case, performing maintenance tasks is the most straightforward mean for keeping the organization on track and thus keeping one's job.

Formally, for survival motive m_s , $expected_satisfaction'(m_s, t, r) = 1$ if $t = t_l$ (the agent obeys the leader) in machine bureaucracies or if $origin(t) = o_m$ (t is a maintenance task) in adhocracies. Otherwise $expected_satisfaction'(m_s, t, r) = 0$.

The satisfaction of conformity motive depends on two conditions related to norms and uncertainty. First, conformity is expected to be satisfied in complying with protocols (standard resolution) for lowering uncertainty. Second, conformity is sensitive to norms, which depends on the organizational structure. In machine bureaucracies, the norm is "legal", consisting of obeying to orders given by leaders. In adhocracies the norm is social, consisting of performing tasks that are the same as those of others.

For conservatism motive m_c , $expected_satisfaction'(m_c, t, r) = 1$ if $r = r_{stand}$ and $t = t_l$ in a machine bureaucracy or else, in adhocracies if $origin(t)$ is the most frequently performed by other agents. Otherwise, $expected_satisfaction'(m_c, t, r) = 0$.

Finally, the openness-to-change motive is exclusively satisfied by creative task resolution.

For openness to change motive:

$$expected_satisfaction'(m_{otc}, t, r) = \begin{cases} 1 & \text{if } r = r_{crea} \\ 0 & \text{otherwise} \end{cases}$$

Generate Plans for New Goals

If a task has been finished last round, agents drop all existing intention : they were related to the previous goal which is either unfeasible or achieved.

When a new goal $solve(t)$ has just been adopted, a new plan is generated for it. This plan consists of working until the end (success or failure) of the current task. This plan consists of three parts: $accept(t)$; $set_resolution_strategy(r)$; $work()$ where:

$$r = \operatorname{argmax}_{r' \in R} \operatorname{expected_satisfaction}(t, r')$$

Select Action

The previous step raises only one plan at a time. Agents perform the first action of this plan.

Conclusion

This section introduces concepts of our model and related formulas for modeling them. As a general description of underlying dynamics, these formulas aim at capturing the main dynamics of how an agent is expected to behave.

Agents aim at satisfying their motives, starting by the survival motive and then higher-level motives. When the organization performs poorly, agents rely on quick-fixes and perform maintenance tasks (adhocracy) or obey to the leader (bureaucracy). This choice is motivated by a high need level and importance for the survival motive and the highest satisfaction provided by the quick-fix resolution which is the most efficient, even if it does not satisfy other motives. When the situation of the organization is stable enough, agents can start to take less efficient resolutions for fulfilling their important motives: acting more creatively or according to norms and standards.

In this setting, culture influences higher-level motives. Thus, agents are more or less satisfied by and so driven by relying on standards or on creativity. If conservatism (high PDI-CD and high UAI-CD) is important, individuals tend to prefer relying on standards and complying with leaders (if any). Conversely, if openness-to-change is important (low PDI-CD, low UAI-CD) individuals are more likely to take more time for acting creatively and thus proposing original solutions (in former versions, potentially disobeying for performing more creative tasks). These behaviors match expectations from culture about individual behaviors as depicted in Chapter 3.

The difference in terms of efficiency and satisfaction between the various resolutions offer rich individual and collective dynamics which are particularly explored in the following section.

5.4 Comparing Our Model with Properties of Human Culture

This section checks whether our model replicates properties of the influence of culture depicted in theories of culture. By showing this replication, we provide further evidence for answering Research Question 2. This replication indicates that values are a key aspect

of the influence of culture on coordination (Research Question 2.a) and our decision architecture illustrates how to model practically this aspect (Research Question 2.b).

In order to perform this check, we need first to consider the core properties that can be explored in our model. We focus on core properties of culture, which are the most validated and general. Then, for each property, we check whether this property occurs in our model.

Which properties can be investigated? Existing theories indicate two main categories of properties. First, properties about the influence of culture on individual decisions. These properties can be explained symbolically directly from our model. Second, properties about the influence of culture on collective outcomes. These properties are better explained experimentally, in running simulations.

This section is structured as follows. First, Section 5.4.1 checks whether our model replicates properties of the influence of culture on individual decisions by using symbolic analysis. Then, Section 5.4.2 introduces our experimental setup and indicates how to process results. Finally, Section 5.4.3 introduces our experiments for checking whether our model replicates properties about the influence of culture on collective outcomes.

5.4.1 Validating the Influence of Culture on Individual Decisions

Existing theories of culture (e.g. Hofstede et al. (2010a), in Section 2.2, Chapter 3) present numerous core properties of the manifestations of the influence of culture on individual decisions. This section investigates whether our model replicates these core properties.

In this chapter, we consider two types of cultural influences. Culture influences the relative importance given to values of conservatism and openness-to-change. The link between these the influence of culture on these values and concrete cultural manifestations is presented in Section 3.2. In particular, in Section 3.2, we introduce empirical and conceptual relationships between the influence of culture on these values and HCDs. In turn, HCDs depict global patterns of manifestations of the influence of culture.

Conservatism: Conservatism is conceptually and empirically related to high PDI-CD and high UAI-CD. These cultural dimensions are related to multiple core behavioral manifestations.

High PDI-CD is related to the core property of high compliance towards instructions provided by leaders. This property is visible in considering our model: conservatism increases the importance of the need for conformity. If a leader is present, when the organization has a bureaucratic structure, then agents obey orders for satisfying their conformity motive. Thus, if conservatism is important, agents are more likely to comply with instructions of leaders.

High UAI-CD is related to the core property of relying on rules and formal interactions for avoiding uncertainty. This property is also visible in considering our model. Conservatism raises the importance towards satisfying needs for conformity. At decision time, these needs are expected to be satisfied by relying on the standardized resolution strategy. Thus, the more conservatism is important for an agent, the more likely this agent will rely on a standardized resolution strategy.

Openness-to-Change: Openness-to-change is related to low PDI-CD and low UAI-CD.

Low PDI-CD is characterized by a relative autonomy towards leadership. This property is difficult to consider in the current model, because leaders propose tasks that are optimal

with regard to the satisfaction of individual desires. This property was more visible in a former version, where leaders could give suboptimal instructions (t_i occasionally pointed to a tasks for which the agent is not an expert at while another task was available). In this setting, if the organization is successful enough (thus survival motive is sufficiently shut off), agents giving more importance to openness-to-change were occasionally disobeying for better satisfying self-expression motives.

Low UAI-CD is related to the core property of relying on adaptive resolutions. This is visible in our simulation: openness-to-change raises the importance towards satisfying the need for self-expression. This need is satisfied in turn in performing creative (or adaptive) resolutions.

In conclusion, the model of culture presented in this chapter and based on our decision architecture replicates the core properties of the influence of culture on individual decisions depicted by theories of culture.

5.4.2 Experimentation Setup

Explorations performed in Section 5.4.3 rely on a few shared input variables. The current section introduces these variables and proposes general explanations for reading and interpreting our plots. These explanations provide a basis for understanding the various plots presented in Section 5.4.3.

Formalizing the Input Space

Our experiments introduced in Section 5.4.3 rely on three core variables: culture, environmental complexity and organizational structure.

Culture: Culture is represented by the relative importance given to values within the value-system and is represented by $\alpha \in [0, 1]$. The higher α the more agents culturally prefer conservatism over openness-to-change.

Environmental Complexity: Environmental complexity represents the inherent complexity of the environment. The higher this complexity is, the more often the organization encounters tasks for which only a few or no agents have the adequate expertise for handling these tasks. Formally, this complexity is represented by $|E|$. The higher is $|E|$, the wider is the range of possible expertise for a task. Since the number of agents remains the same, the higher $|E|$ is; the lower the probability that the expertise of an agent matches the requirements of the available task.

Organizational Structure: We consider two organizational structures: machine bureaucracies and adhocracies. The difference between these two patterns on the simulation are described in Section 5.1 and formalized in Section 5.2.2.

Filling and Reading Output Plots

We explore a four-dimensional space: three dimensions are required for the input and at least one dimension is required for representing the output. This number of dimensions is not conveniently laid on the 2-dimensional pages available for printing this thesis.

In order to handle this high dimensionality, we remove one input dimension and represent it through independent plots. In more details, we remove the “organizational struc-

ture” variable which has only two values. Thus, each plot is entirely either “in a machine bureaucracy” or “in an adhocracy”.

The three remaining dimensions can be handled with colored grids. The position of each cell expressed in terms of (X, Y) coordinates. These coordinates indicate the values for the two input variables. More precisely, X axis indicates the value of $\alpha \in [0, 1]$ and Y axis indicates the value of $|E| \in [1, 40] \cap \mathbb{N}$. The color of the cell indicates the value of the output variable: the darker, the higher this value is. This variable is later defined depending on the experiment. The output value of each cell is based on the average output of 20 runs. Each run lasts for 150 rounds each with 15 agents. These values were sufficient for obtaining stabilized outcomes.

Here are some general guidelines for reading figures. In addition, the more a cell is on the left, the lower is α , the more openness to change is important with regard to conservatism. The reverse applies when going to the right. The more a cell is close to the bottom of the figure, the lower is $|E|$, the simpler is the environment.

In order to get a more general overview of the correlations depicted by a figure, a comparative study is required. In considering a single row, the environmental complexity is fixed and only culture varies. This can be used in particular for observing gradients (e.g. the output value increases when α decreases). These gradients indicate a correlation between the output value and the variable of culture. Similar observations can be done in considering the graph vertically, for considering the influence of the environment on the outcomes.

5.4.3 Validating the Influence of Culture on Collective Outcomes through Experiments

This section introduces basic properties of the influence of culture on collective outcomes as depicted in theories of culture and checks, through experiments, whether these properties are replicated by our model. In particular, given the end-goal of this thesis, we focus on the influence of culture on performance.

Before investigating properties of culture, we first need to check whether our model reproduces basic properties of the influence of type of organization on collective performance. Indeed, we consider two organizational structures that are known to support two very different performance profiles in human societies, independently from culture. In order to avoid confusing the influence of the type of organizations on the results with the influence of culture on the results, we first check whether our model of organization is coherent with expectations from reality.

Then, we investigate relationships between culture and collective outcomes (or organizational performance). Culture is known to influence performance *relatively independently* of implemented coordination mechanisms. This relative independence raises two types of properties: generic properties qualifying a direct influence of culture on collective outcomes and more specific properties considering the influence of culture on collective outcomes with the partial influence of coordination mechanisms. These later properties were actually the core initial target of this model: as highlighted in Figure 2.3, conservatism is known to match well with machine bureaucracies while openness-to-change matches well with adhocracies.

Some core direct links between culture and performance are studied by Hofstede et al.

(2010a). First, UAI-CD and PDI-CD are known to be correlated to efficiency in simple environments. Thus, conservative cultures should achieve high efficiency in simpler environments. Second, UAI-CD is known to be negatively correlated to organizational flexibility (the lower the UAI-CD, the more flexible). Thus, open-to-change cultures should display high flexibility.

Some other links, taking into consideration the organizational structure, expand these direct links. Gibson and McDaniel (2010); Hofstede et al. (2010a) suggest that different types of culture lead to better individual satisfactions and collective performances depending on the underlying coordination structure. A general illustration of these theories is presented in Figure 2.3. In particular, these theories indicate that machine bureaucracies match high PDI-CD and high low UAI-CD (conservatism in our model) while adhocracies appear to match low PDI-CD and UAI-CD (openness-to -change in our model).

Following sections introduce each experiment for studying whether our model replicates these properties.

Generic Culture-Independent Organizational Properties

Before considering how cultures influence agents and collective performance, we want to check first that our model of organization alone replicates basic properties as observed in human societies. In more details, we want to check whether our model replicates the influence of the organizational structure on performance, as suggested by Mintzberg (1979). This theory explains that machine bureaucracies tend to be more efficient than adhocracies for handling simpler environments and adhocracies are more adequate for handling more complex environments.

In order to study whether this property is replicated by our model, we run several simulations. More precisely, we investigate the differences in terms of efficiency between bureaucracies and adhocracies for a given cultural and environmental setup. If the environment is simple enough, bureaucracies should achieve higher performance, if the environment is more complex, adhocracies should achieve higher performance. This property is culturally-independent.

This property is tested in Figure 5.5 (this Figure is based on the merge between the two plots from Figure 5.6). Each cell of this figure introduces the difference in terms of efficiency between adhocracy and machine bureaucracies. Formally, efficiency is measured through the OS variable. The greyscale of the cell (X, Y) measures

$$OS_a - OS_b$$

where OS_a is the OS value of the adhocracy and OS_b is the OS value of the bureaucracy for which parameter α (indicating culture) is set to X and environmental complexity $|E|$ is set to Y .

The property is verified. Independently from the cultural setting (for any column of the graph), bureaucracies achieve better results than adhocracies (darker values) when environmental complexity is low (bottom of the graph). When the environmental complexity increases (moving from bottom to the top of the graph), the more often adhocracies achieve higher performance than bureaucracies and the larger is this difference (the lighter is the cell). This property can be considered from an overview in observing a gradient along Y -axis.

Additional explanations can be provided by considering our model. In machine bureaucracies, leaders try to allocate tasks for which the subordinate is an expert without considering the task type. In adhocracies, agents try to solve maintenance tasks first for satisfying their survival motives. Thus, they may face the dilemma “solving a task quickly vs. solving a maintenance task satisfying survival drives”. In simpler environments, the best solution consists in solving tasks efficiently for achieving the highest performance through many successes (machine bureaucracy). In more complex environments, where expertise is harder to get, it is better to solve maintenance tasks which are the most important for the organization (adhocracy).

Directly Relating Culture and Performance

Hofstede et al. (2010a) relate cultures to collective performance. In particular, the PDI-CD is related to good performance in simple environments and low performance in complex environments. The UAI-CD is related to high efficiency in simpler environments but lower flexibility. In order to check that in our model, conservatism should support high efficiency in simpler environment-and-openness to change should support better flexibility and good efficiency in complex environments.

Efficiency: In order to check whether our model replicates properties of the influence of culture on efficiency, we measure the efficiency achieved by the organization depending on the organizational structure, culture and environmental setup. These measures are represented by Figure 5.6. Efficiency is measured by the *OS* variable.

Our model verifies the relationship between conservatism and high efficiency in simpler environments. This is particularly visible in machine bureaucracies. This can be observed in considering the first lower row(s) of the graph. These rows represent the simplest environments. In these rows, the higher X is (or α or preference towards conservatism), the brighter is the cell (the higher is the performance).

To some extent, this observation is also visible with adhocracies (which matches poorly with conformism as studied in the following section). This can be seen in considering the first few rows. In each row, a “front” of whiter cells (indicating good performance) can be seen. Initially this front is in the right side (around 0.8), indicating a good match between

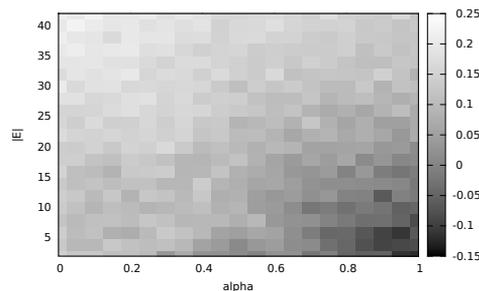


Figure 5.5 – Comparing efficiency between adhocracies and bureaucracies depending on α representing the relative preference towards openness-to-change (left) vs. conservatism (right) and environmental complexity (bottom to top). Lighter tiles, (with a value above 0) indicate that adhocracies perform better than bureaucracies in this culture and this environment.

conformism and good performance. In increasing this environmental complexity (when reading one row up), the more this “front” is pushed towards the left (the more openness-to-change is related to better results).

The relationship between openness to change and high efficiency in more complex environments is visible with adhocracies. In considering the upper rows, lighter cells (indicating good efficiency) are on the left side of the figure (indicating a cultural preference for openness-to-change). This link is not visible with bureaucracies due to the too strong mismatch between openness-to-change and bureaucracy.

Flexibility: In order to check whether our model replicates properties of the influence of culture on flexibility, we measure the flexibility achieved by the organization depending on the organizational structure, culture and environmental setup. Flexibility indicates the capability for a system for maintaining good performance in spite of changing environments. In order to measure flexibility, we consider the degree which efficiency decreases when the environmental complexity increases. Formally, flexibility is modeled by the comparison between OS_n and OS_{n+1} where OS_i is the organizational success (and thus efficiency) when $|E| = i$. In simpler terms, the darker is the cell, the more the organization loses in terms of performance when increasing complexity, and the lower is organizational flexibility.

Variations in terms of flexibility are presented in Figure 5.7. The property is verified for both organizational structures: when environmental complexity increases, organizations with cultures promoting conservatism suffer from to highest performance losses compared to those promoting openness-to-change.

This property can be observed in Figure 5.7 *a* and *b*. For both *a* and *b*, darkest zones (highest performance drop), are more frequent, more important and happen earlier when conservatism is the most important (right side), than when openness-to-change is important (left side). Therefore, performance in conservative culture tends to be less flexible than performance with more open to change culture. Since this phenomenon happens in both figures, then culture influences the collective flexibility independently of a given organizational structure.

This loss of flexibility can be explained in terms of our model. When environmental complexity is low enough, agents often have the adequate expertise for handling tasks.

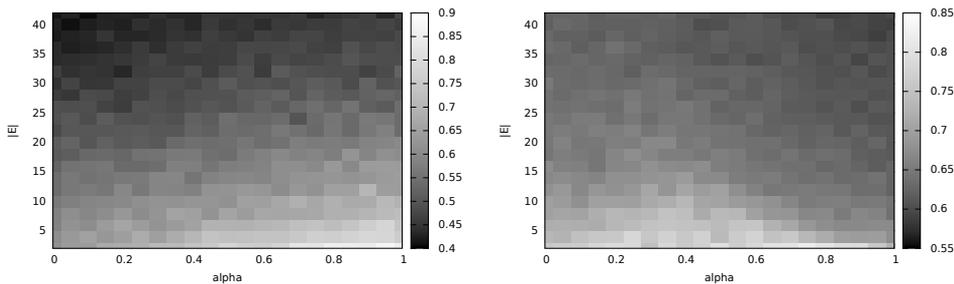


Figure 5.6 – a, (left) and b (right) Efficiency in bureaucracies (left figure) and adhocracies (right figure), depending on the relative cultural importance (X-axis) towards creativity (left side of each figure) vs. conformity (right side of each figure) and environmental complexity (Y-axis). The lighter is the tile, the higher is the OS of the organization

Thus, agents can trigger procedural resolutions (supported by conservatism) more easily and more efficiently than creative resolutions (supported by openness-to-change). Nevertheless, when complexity increases, procedural resolutions become much more expensive to trigger than creative resolutions: attempting procedural resolutions (supported by conservatism) lead to more failures than attempting creative resolutions (supported by openness-to-change).

Relating Culture, Coordination Mechanisms and Performance

In addition to direct general influences of culture on collective outcomes, the influence of culture on collective outcomes is also known to be sensitive to coordination mechanisms. More precisely, pairs of the form (culture, coordination mechanisms) are known to match each other well and thus to lead generally to higher performance. In this direction, we introduce several mechanisms for explaining this match in Chapter 6.

In the case of our model, Figure 2.3 (and underlying theory from Hofstede et al. (2010a)) depicts that conservatism matches machine bureaucracies and openness-to-change matches adhocracies. Two matches can be considered in our model: efficiency and satisfaction.

Culture, Coordination Mechanism, Efficiency: In order to study this property, we need to consider whether our model matches four properties: (1) conservatism leads to highest performance in bureaucracies (2) openness-to-change leads to the highest performance in adhocracies, (3) conservatism leads to worse performance in adhocracies and (4) openness-to-change leads to worse performance in bureaucracies. Given our experimental setting, verifying (1) and (2) entails (3) and (4).

In order to check the property (1), we consider whether cultures supporting conservatism tends to achieve higher efficiency than culture supporting openness to change in bureaucracies. This property can be verified in considering Figure 5.6. For all environmental setups (in setting $|E|$ and thus reading the graph along a horizontal “row”), a gradient can be observed. The more the cell is on the right (the higher α , the more conservatism is culturally important), the lighter is the cell (the more efficient is the organization).

In order to check property (2), we consider whether culture supporting openness-to-

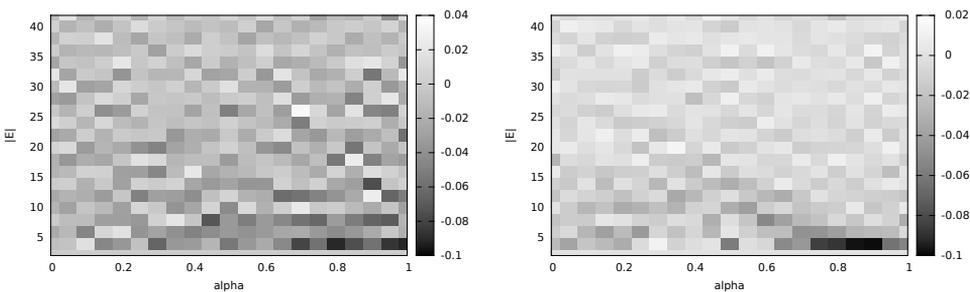


Figure 5.7 – a (left) and b (right) Flexibility (maintenance of good performance when increasing environmental complexity), depending on culture (X-axis) and on environmental complexity (Y-axis) in bureaucracies (left) and adhocracies (right). Lighter tiles indicate that the performance does not decrease in spite of increasing environmental complexity

change achieves higher efficiency than cultures supporting conservatism in adhocracies. The procedure is the same as for the first property, but in observing a reverse gradient. The more the cell is on the left (the lower α , the more openness to change is culturally important), the lighter is the cell (the more efficient is the organization). The only exception to this rule concerns very simple environments. In that case, agents have almost always the required expertise. Thus, standardization, supported by conservatism leads to higher efficiency. Nevertheless, this tendency fades away when environmental complexity increases.

Culture, Coordination Mechanism, Satisfaction: Gibson and McDaniel (2010); Hofstede et al. (2010a) suggest that satisfaction of individuals depends on the combination between culture and coordination mechanisms.

In order to investigate this property, we evaluated the average satisfaction of agents at the end of a run. This value is the average satisfaction of agents, where satisfaction is represented by:

$$satisfaction_a = \sum_{m \in M} importance(m) \times sat_a(m)$$

Results of this experiments are displayed in Figure 5.8, which relates individual satisfaction with culture, organizational structure and environmental complexity. Lighter cells represent higher satisfaction.

With a similar analysis as for the case of efficiency, Figures 5.8 *a* and *b* display a horizontal polarization; (1) in bureaucracies, satisfaction is correlated to conservatism while (2) in adhocracies, satisfaction is correlated to openness-to-change. Once again, a correlation between satisfaction, organizational pattern and culture is observed.

These properties can be further explained conceptually in considering our model. For property (1), decisions supported by conservatism match decisions supported by the coordination mechanism. Thus, agents can regularly satisfy both their survival and conformity motives. Conversely, creative resolutions imply working more slowly, forcing to make a choice between satisfying survival and self-expression motives. For property (2), decisions supported by openness-to-change match decisions supported by the coordination mechanism. Indeed, for satisfying their survival motives, agents can only access a limited range of tasks (maintenance tasks). This limitation can prevent agents to apply procedural resolutions (cannot access the adequate task, resolution is too expensive). Conversely, in

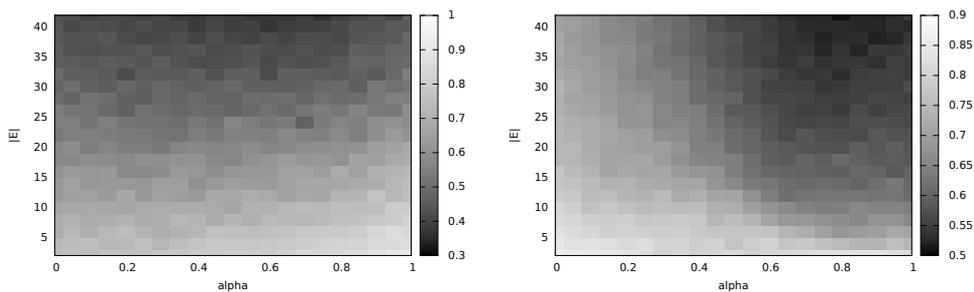


Figure 5.8 – a (left) and b (right) Individual satisfaction in bureaucracies (left) and adhocracies (right) depending on culture (X-axis) and environmental complexity (Y-axis). Lighter tiles indicate that individuals are more satisfied in this organizational structure, environmental complexity and cultural setup.

complex environments, creative resolutions can be applied with much more limited costs and thus be triggered more easily.

Note the particular case of low environmental complexity and a cultural preference towards openness-to-change preference, in combination with the previous experiment: individuals are more satisfied in spite of lower efficiency (so, lower satisfaction provided by survival drives). This case corresponds to organizations confronted to a friendly environment. Since survival motive is easily satisfied, agents can afford to make decisions for satisfying motives supported by their values at the expense of efficiency.

Exploring Beyond Theories: Matching Culture, Organization and Environment for Achieving High Performance

Theories such as Hofstede et al. (2010a) implicitly suggest that a match between culture, organization and environment is important for achieving high performance. Studies presented in this section highlight the importance of this match. In particular, this section highlights two winning combinations: conservatism, machine bureaucracy in simple environments and openness-to-change, adhocracies in complex environments.

All these aspects are visible in Figure 5.5. The brightest cells, where adhocracies are the most efficient in comparison with bureaucracies occur on the top-left corner. In this corner, culture is the most open to change and the environment is the most complex. Conversely, the darkest tiles, where bureaucracies are the most efficient in comparison with adhocracies occur in the bottom-right corner, when culture supports conservatism the most and the environment is the simplest.

This observation highlights that these aspects are to be taken into consideration when using culture for supporting coordination. Furthermore, this observation suggests that there is no single culture which dominates others.

Conclusion

This section shows evidence, through symbolic and empirical analysis, that our model replicates core properties of the influence of culture on coordination, as observed in social sciences. These numerous relations provide more confidence that our model, and the architecture on which this model is built on, replicates the core properties of the influence of culture.

Additionally, these experiments replicate the tight match between culture, coordination, environment and collective performance. This match is at the core of the motivation of this thesis, since it suggests that taking culture into consideration is crucial for improving coordination. The fact that we managed to reproduce it with our model strongly suggests that our model can be used for practically improving coordination in artificial societies.

Furthermore, this experiment matches a core suggestion from theories: no culture “dominates” all the others. Instead, a specific culture is adequate for different environments and should be streamlined with the adequate organization for achieving the highest performance.

5.5 Additional Evidence

Through this chapter, we have built a simulation and tested it against properties of culture provided by theories of culture. The match between properties of culture from our model and those of real culture provides some good confidence that: (1) values are a key aspect for considering the influence of culture on coordination, and (2) our decision architecture from Chapter 4 is appropriate for replicating the basic influence of culture on coordination within artificial societies.

Nevertheless, further confidence can be required for considering the validity of our decision architecture (e.g. would it work with another set of values or in a different organization?). This section provides additional evidence for acquiring this confidence.

The simulation model and its extensive match the numerous properties of culture depicted by theories provide a first piece of evidence that our decision architecture makes sense. In particular, explanations we provide for explaining collective phenomena are sound with regard to theories of culture.

As a second piece of evidence, we expanded our simulation with additional values and obtained results that match reality. In order to do that, we added two values also from Schwartz (1999): achievement and benevolence (MAS-CD-related values). In more details, the organization was providing rewards to agents and agents were given the possibility to work together. Those preferring achievement were more driven by rewards and thus more likely to handle tasks by themselves. Those preferring benevolence were more likely to work together. Our simulation lead to similar cultural properties as presented in Hofstede et al. (2010a) (high MAS-CD is related to higher performance in simple environments, low MAS-CD is related to higher performance in complex environments). When achievement was important, individuals perform better in simpler environments (agents are experts, they can handle tasks alone before reaching the deadline, avoiding costs for initiating the cooperation). Conversely, when benevolence was important, organizations were more adequate for managing more complex environments (cooperating was paying off better for handling complex tasks in due time).

As a third evidence, we can refer to Chapter 6. This chapter introduces another simulation based on our decision architecture, with slightly different coordination mechanisms and values. Results from this chapter also match influences of culture which are recognized in social sciences.

As a last piece of evidence, our decision architecture is based on an extensive theoretical backup. The importance of values on the influence of culture is introduced by Hofstede et al. (2010a). Miceli and Castelfranchi (1989) further shows the importance of values for supporting social behavior. Furthermore, our decision architecture replicates some core properties of values depicted by Miceli and Castelfranchi (1989) for supporting coordination: we make sure that agents decisions are in line with their values.

All this evidence provides good confidence that values have a core impact on the influence of culture on decisions and that our decision architecture is adequate for reproducing the influence of a wide range of values and that this influence impacts coordination as in human societies.

5.6 Chapter Summary

This chapter provides evidence that our architecture introduced in Chapter 4 can be used for building decision models that replicate core properties of human culture. In more details, our decision architecture, can be used for building models that replicate the influence of culture on individual decisions, but also on collective outcomes.

While we do not claim that our architecture is “the only culturally-sensitive architecture”: more architectures can be provided. However, the validation introduced in this chapter shows that our architecture is relatively accurate for building models of the influence of culture that match the influence of culture as observed in human societies. This validation gives sufficient confidence for moving on: by using it carefully, we can realistically expect to replicate human-like behaviors.

Regarding following chapters, the confidence gained in the current chapter is used for better understanding the influence of human culture (Chapter 6) and for studying how to use models of human-like culture for supporting coordination (Chapter 7).

Answering Research Question 1: *How Does Culture Influence Coordination in Human Societies?*

This section highlights an apparent contradiction between two properties of the influence of culture on collective outcomes in situation of coordination. Theory suggests that culture has an influence on collective that is generic and relatively context-independent (e.g. high PDI-CD achieves better results in simpler environments). However, the influence of culture on collective performance is also considered to be dependent on coordination (e.g. high PDI-CD culture does not match adhocracy). Sometimes, this coordination-sensitive influence can take over the generic influence (e.g. in our simulations, high PDI-CD cultures achieves lower performance in adhocracies than low PDI-CD cultures).

Chapter 6 proposes an explanation for solving this apparent contradiction.

Answering Research Question 2: *How to Model The Influence of Culture on Coordination?*

This chapter reinforces the answer to this question provided in Chapter 4 by showing evidence for supporting this answer. We modeled the influence of culture on coordination through values. We obtained results that match expectations from social sciences, at individual and collective levels.

Answering Research Question 3: *How Can We Use Models of Human Culture as a Practical Tool to Support Coordination in Artificial Societies?*

When considering using culture for supporting coordination, this chapter suggests a very important rule to take into consideration: seemingly, there is no single “better” culture that dominates all other cultures. Thus, when using culture for supporting coordination, system designers should consider the match between culture, environment and other coordination mechanisms.



6

Explaining the Influence of Culture on Coordination from Individual Decisions to Collective Outcomes, Theory and Simulations

Chapter Summary

This chapter proposes a theory for better explaining how culture influences coordination. More specifically, this theory explains how the known influence of culture on individual decisions leads to known influence of culture on collective outcomes. This theory is then illustrated by several examples arising from simulation models.

Available theories of culture explore several relationships between culture and coordination. More precisely, these theories explore mainly two types of links, as illustrated in Figure 6.1: theories explore the (rather direct) influence of culture on individual decisions (e.g. individuals with high PDI-CD culture tend to be more compliant towards orders); and the (rather indirect) influence of culture on collective outcomes (e.g. high PDI-CD organizations tend to perform better in simple environments).

However, these theories do not explain how known influences of culture on individual decisions lead to known influences of culture on collective outcomes. Understanding this link is important for achieving the end-goal of this thesis. Indeed, as MASs designers, we need to have at least a basic understanding about how the influence of culture on collective performance results from the influence of culture on individual decisions. This chapter particularly aims at answering Research Question 1.b: *How does culture (indirectly) influence collective outcomes related to coordination in human societies?*

This chapter aims at investigating this link through a theory. This theory inspired by concrete evidence provided by existing theories of culture, mostly Hofstede et al. (2010b) as presented in Section 2.2 and expanded in Chapter 3. In addition, the theory presented in the current chapter is based on our own experiments conducted through simulations. These simulations rely on our model of the influence of culture on individual decisions,

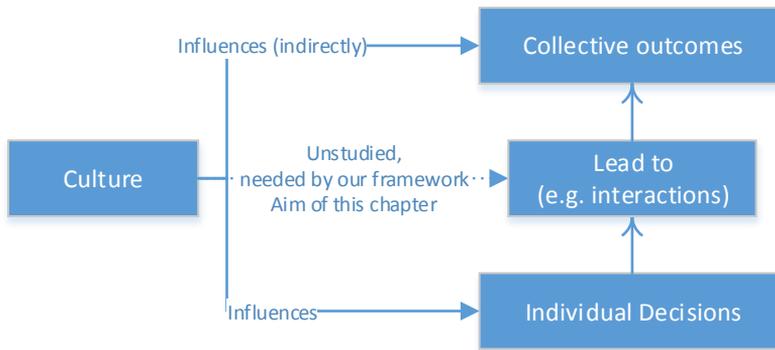


Figure 6.1 – Existing studied links relating culture, individual decisions and collective outcomes. In this chapter, we aim at filling the link in the middle, allowing to relate how the influence of culture on individual decisions leads to the (known) influence of culture on collective outcomes.

introduced in Chapter 4 and that is shown to match with expectations of the influence of culture on individual decisions and collective outcomes, as depicted by social sciences in Chapter 5. This theory introduces in particular several concepts that we can be used for explaining the influence of culture on coordination. These concepts will be used in Chapter 7 as a basis for determining how cultures can be concretely used for supporting coordination.

In order to illustrate how this theory applies in concrete situations, we present the experiments that we designed in order to get inspiration for creating this theory. This technique is described in details in Epstein (1999). We acknowledge that a complete investigation would require an extensive empirical cross-cultural study, in order to provide exact empirical evidence. Such a study is far beyond the scope of this thesis, which aims at exploring a new design technique. Instead, simulations offer a simple yet rich, stable, manipulable and explorable playground in which we could easily test our hypotheses and which now serves as a concrete illustration. Simulations allow to easily try out a variety of coordination mechanisms and of types of cultures. Furthermore, findings from these simulations are sufficient for providing credible explanations. While not necessarily being exactly conform with reality, this explanation is close enough for making a parallel with artificial systems and being operationally useful. Finally, we do not ignore the realism of our simulations by comparing them with empirical findings from social sciences.

This chapter is structured as follows. First, we introduce our theory in Section 6.1. Then, using this theory as a basis, we introduce simulations for illustrating it in concrete scenarios. Section 6.2 introduces an informal overview of the simulation domain. This overview is formalized into an abstract simulation model in Section 6.3. Then, following sections are concrete implementations of this abstract simulation model.

6.1 A Theory of the Influence of Culture on Coordination

Current available theories of culture are limited when explaining this link. These theories present numerous examples of the influence of culture on individual decisions and of the (indirect) influence of culture on collective outcomes in situation of coordination. However, while giving correlations between individual and collective levels, they do not introduce the *general mechanisms* for explaining how cultures, individual decisions, collective outcomes and coordination relate to each other. This section aims at proposing a theory for further connecting these various perspectives and for introducing these general mechanisms.

Our theory relies on basic concepts that should be described first for framing this theory. Our theory is compliant with the MAS framework, as we consider an agent-based setting. A group of individuals or agents is embedded within an *environment*. Each agent locally perceives the environment and makes decisions according to its own *decision process*. In our context, we assume that agent decision processes are influenced by *coordination mechanisms* (e.g. norms, organization, goals set by a coordinator) and by a *culture*. In particular, we focus on *values* as introduced in previous sections. Joint decisions of agents lead to *collective outcomes*. Given this framework, our aim consists in explaining how (known) influences of culture on individual decisions lead to (known) influences of culture on collective outcomes, depending on the environment and the influence of coordination mechanisms.

As a first attempt to explain the individual-to-collective link, one may try a direct connection between the two. Nevertheless, our attempts and those from available theories in this direction show that this connection is hard to make directly. In general, describing this connection requires introducing several elements specific to the given cultural phenomena (e.g. in high PDI-CD culture, status is more important, leaders are given more responsibilities). Introducing these specific elements make that explanations are *ad-hoc* (e.g. the explanation for PDI-CD can hardly be translated for explaining UAI-CD). As a consequence, while having explanations for specific phenomena, these direct explanations fail to provide general mechanisms for explaining the influence of culture¹. We need additional concepts and perspectives for explaining this link.

In investigating the patterns used by available theories for explaining the individual-to-collective link, a lot of help can be found in the recurrent perspective of *interactions*. Interactions provide a practical middle ground for bridging the gap between individual decisions and collective outcomes. Interactions are close enough to individuals for being directly connected to individual decisions. Interactions are also easier to relate to collective outcomes than individual decisions. Consider the following example. From an individual perspective, the preference towards obedience supports a preference from both subordinates and leaders to rely on leaders for making collective decisions. From the perspective of interactions, these individual decisions lead to higher centralization. From the perspective of collective outcomes, these interactions lead to higher performance in simpler environments (tight management of subordinates by leaders who can grasp the whole

¹And we need these mechanisms for building realistic and coordination-supportive models of culture in Chapter 7

environment).

Culture influences individual decisions, which influences in turn interactions. Since all individuals are (expected to) share the same culture and they thus share similar influences. Then, when gathering these influences, culture influence interactions.

We can understand in more details how culture (indirectly) influences interactions in further considering the specificities of the influence of culture on individual decisions and mechanisms of culture. These specificities are introduced in more detail in Section 2.2, Chapter 3 and Section 4.1.2. We remind here the core aspects for building our theory. Values are *irrational heuristics*. They indicate whether a decision is “good” according to some perspective (e.g. honesty, obedience, autonomy). This heuristic is not necessarily rational: at best, values indicate aspects that generally lead to individual satisfaction, without guaranteeing that this decision is rationally the best. Values supported by culture can even point to values that are clearly not “individually good” (e.g. self-sacrifice, charity). Nevertheless, these values are generally expected to be *collectively “good”*. Values supported by culture tend to be relatively *abstract*². These abstract values can be implemented concretely in a wide range of decisions, coordination mechanisms or more specific values. Values supported by culture are deeply expected to be *shared*. To a lesser extent, the same applies for their concrete manifestations (e.g. rules, rituals, symbols). Thus, individuals generally create *expectations* about other individuals, the society and the environment based on their values. Sometimes, the core reason that makes that these values actually indicate “good” individual or collective behaviors is because they are shared (e.g. giving orders when expecting obedience as a leader is likely to fail if subordinates are not obedient).

The core of our theory is the following. *Culture, particularly through values, tends to support the occurrence of specific types of interaction patterns. These patterns are abstract, coherent and recurrent.*

Abstract in a similar sense as the object-oriented paradigm: the interaction pattern supported by a culture can be instantiated in concrete situations of interaction (e.g. the value of obedience supports a “leaders-decide-for-subordinates” relationship that can be implemented in many ways depending on the environment: which orders to tell, telling them through meetings or emails and so on).

Coherence is more subtle to define but can be considered as a common drive towards achieving (intentionally or not) the culturally-driven interaction pattern shared by all individuals. From an individual perspective, individuals are driven by values towards making decisions leading to the interaction. This influence can be intentional (e.g. as a subordinate, try to reach a situation where leaders are informed) or not (e.g. tell information when asked by the leader). This notion of coherence can also be considered from an external perspective. This can be seen as individuals pursuing a common (self or collective) interest.

Recurrent means that such an interaction pattern tends to occur relatively independently of a specific interaction situation. Thus, in various situations of interaction, one can expect to see a similar abstract interaction patterns occurring if the culture is the same (e.g. with the value of obedience, see leaders giving orders and subordinates reporting information)

²We focus on these values in this thesis

The influence of culture is also dependent on coordination mechanisms and environments. Coordination mechanisms and environments can support the influence of culture by raising culturally-sensitive decision aspects (e.g. PDI-CD is related to “identity” and it has more expressiveness when the organization structure includes leaders and subordinates). Conversely, coordination mechanisms and environments can also restrict the influence of culture when individuals act rationally (e.g. culture can hardly be expressed in factory work where workers have the mission “push a button when an item comes”). In other words, in assuming that individuals are rational, culture tends to be expressed only in situations that are not well handled by coordination mechanisms or not constrained by the environment (e.g. unexpected situation, flexible coordination, non-obvious dilemma). Culture is particularly visible when individuals have sufficient freedom for making decisions. In this context, culture drives the occurrence of abstract interaction pattern which complements the freedom left out by formal coordination mechanisms and rationality.

In considering culture from a qualitative perspective, culture can be *more or less shared*. The more individuals share a culture (e.g. because they have the same origin), the more culture supports a precise abstract interaction pattern. For instance, values of obedience and security, when culturally-shared, support interactions where leaders manage subordinates through rules as depicted in Figure 2.3. This culture supports more accurate interactions than a culture supporting just obedience, for that any situation where leaders make decisions for subordinates would be appropriate. Furthermore, the more culture is shared the more expectations individuals can create about each other and leading to more coherent interaction patterns. Furthermore, richer cultural backgrounds support greater cultural influences on individual decisions, interactions and thus collective outcomes in a wider range of environments and coordination setups.

In spite of its apparent benefits (raising coherent interactions), sharing a culture does not necessarily lead to better collective outcomes. Indeed, incoherent interaction patterns raised by cultural mismatches can sometimes lead “by accident” to better performance than coherent ones. Furthermore, culturally-driven coherent interaction patterns can also achieve only poor performance due to inadequate coordination and environment settings. Nevertheless, coherence and recurrence supported by shared cultures appears to lead to better collective outcomes in the general case.

The rest of this chapter is composed of simulations which illustrate and underpin the various relationships presented in this theory. From a methodological perspective, we elaborated this theory by using “generative social sciences” as depicted by Epstein (1999). These simulations, in addition to insight from Hofstede et al. (2010a), provided the basic inspiration for elaborating this theory. Multiple relationships presented in this theory actually result from the fruitful confrontation between our simulation models and our initial insight: surprising simulation outcomes helped us sharpening our ideas.

These simulations are presented in the following sections. Section 6.4 introduces a baseline, illustrating the abstract interaction pattern resulting from a culture and showing its coherence. Section 6.7 illustrates the influence of coordination and environments on cultural expressiveness. Section 6.8 illustrates the recurrence by exploring whether similar interaction pattern occur in spite of different coordination structures and environments. Section 6.5 illustrates incoherence resulting from unshared culture as well as the influence of sharing a culture on collective performance. Finally, Section 6.6 illustrates the influence of the richness of the expected cultural background on individual decisions coherence and

collective outcomes.

6.2 Overview

This section introduces an overview of the abstract simulation scenario that we explore throughout this chapter. A formal model of this overview is introduced in Section 6.3.

As an overall remark, for sake of continuity throughout our simulations, this simulation model is very similar to the one presented in Chapter 5. This simulation model also consists of an IT support organization trying to help its clients, receiving and processing tasks. The source of the differences separating the two models is shift of our focus between their use. In the previous simulation model, interactions were kept as simple as possible (agents were sharing the same task pool). In the current model, we aim at raising richer and more direct interactions. Thus, environments and coordination mechanisms are altered such that they enforce more interaction-oriented decisions.

Section 6.2.1 introduces the environment, Section 6.2.2 introduces the coordination model and Section 6.2.3 introduces the core aspects of agent decision processes.

6.2.1 Environment

The environment consists of incoming requests made by clients. These requests are processed as *tasks* by the organization. These tasks are given an explicit *deadline* before which they have to be solved. Each task requires a certain amount of work in order to be solved, which can be translated into a *duration* because agents are assumed to work at the same speed.

Sometimes, clients can raise *complex tasks*. These tasks correspond to tasks which require more time than expected in order to be solved. As an example of a complex task, a client calls with the problem “my email does not work”. After solving a problem with the mail server, the organization discovers that the network is also faulty and must be fixed in order to solve the problem. In that case, the *deadline* is expanded in order to handle the extra *duration* raised by this complexity.

Throughout the simulations performed in this chapter, the organization will face a variety of environments. Each environment has some characteristics which pushes the organization to a specific limit. The environment can be *intensive*, with steadily incoming tasks which have to be quickly solved. This environment is used for testing organizational capability for efficiently allocating tasks. The environment can be *complex*, with a relatively slow input of incoming tasks but which have a chance of being complex. Tasks of this environment, in spite of possibly taking more time than expected, remain relatively easy to handle by the organization (long deadline, enough processing power). This environment is used for testing the influence of environmental complexity on the organization. Finally, the environment can be *irregular*, in which many tasks arise at once. This environment is used for studying organizational robustness to congestion.

This environment introduces the challenge of managing the workload of agents, in the presence of potential uncertainty caused by partially observable task durations (for complex tasks). This environment is appropriate for raising interactions in a leader-subordinate organization setting.

6.2.2 Coordination Model

The organization is referred to as DO-IT-2.0, the reorganization of the Department Of Information Technology presented in Chapter 5. DO-IT-2.0 is based on a simple structure: direct delegation from leaders to subordinates.

As for DO-IT, new tasks are issued when receiving calls from clients. However, for DO-IT-2.0, these tasks are only accessible by the leader. Leaders are in charge of allocating tasks to subordinates. Subordinates are in charge of determining whether they accept allocated tasks and of working on accepted tasks until their completion (success or failure). When subordinates have accepted to perform multiple tasks at a given time, they must perform them in a “first-in, first-out” order.

This organization is used in the baseline simulation model. It is referred to DO-IT-2.0. Throughout our simulations, we explore some variants of this coordination mechanism. Rule-DO-IT-2.0 expands DO-IT-2.0 by adding strict regulations for telling how agents should behave: leaders must allocate tasks by using a round-robin and subordinates must accept them. A third organization, referred to as DO-IT-2.0-Together is an open variant of DO-IT-2.0, supporting free cooperation between subordinates and allowing them to share tasks.

6.2.3 Agent Decision Processes

The environment and the coordination model support core decisions that agents have to make. These decisions are in particular related to two core individual decision aspects: identity and uncertainty. After having introduced these decisions and decision aspects, we relate them to values that are at play in this model.

Decisions

The environment and coordination model raise several decisions that agents have to perform. Leaders have to decide how to allocate tasks: shall they take time and monitor subordinates or allocate them more quickly? They have to determine whether they allocate a task and if they do, whom they allocate the tasks to. Subordinates have then to decide whether they accept allocated tasks.

Decision Aspects

Decisions to be performed by agents can be related to two main decision aspects considered in Section 4.4.

The first aspect is *identity*: DO-IT-2.0 features the roles of leaders and subordinates which are strongly differentiated in terms of identity (role, duties, interactions). Furthermore, interactions occur between agents with different roles.

The second aspect is *uncertainty*: The environment tackled by DO-IT-2.0 features uncertainty caused by partial information. Agents cannot know for sure whether a task is complex or not and whether more time will be required for processing it. Thus, subordinates cannot have any certainty about the duration of their remaining work. Same way, leaders cannot exactly estimate the workload of their subordinates.

Culture

Our model of culture consists of two opposed values: autonomy vs. obedience. Another pair of values, assertiveness vs. modesty is explored in Section 6.6.

Autonomy and obedience are conceptually strongly related to the two extremes of the PDI-CD, as introduced in Section 3.2. In our simulations, these two values directly influence the impact of the “identity” aspect on decisions.

Autonomy can be translated by the preference that each individual can and should manage itself, without the involvement of a leader. Autonomous individuals are more likely to prefer avoiding wasting time exchanging personal information. For agents that give high importance to autonomy, leaders should propose tasks and let subordinates free to decide whether they accept it or not. In other words, subordinates should manage their own workload.

Obedience, or social order or hierarchy, can be translated by the preference that dedicated individuals (here, leaders) are in charge of managing other individuals. Obedient individuals are more likely to accept to reveal personal information to leaders for the sake of coordination and leaders to expect to access to this information. In a culture where obedience is important, leaders should tightly manage subordinates and subordinates should comply with instructions from leaders.

Conclusion

The coordination problem introduced in this section densely interrelates culture, coordination, environment, decisions and interactions between agents. Indeed, this model introduces an environment that raises the problem of optimizing task allocation to a group of agents in a situation of uncertainty (tasks can take longer than expected in order to be processed). An organization, based on the structure of the “simple structure” from Mintzberg (1980) tackles this environment by letting an agent, the leader, allocating tasks to subordinates who have the role of processing them. This environment and coordination mechanisms are tightly related to two core decision aspects (identity and uncertainty) which are themselves particularly well connected to the values we investigate (autonomy and obedience and assertiveness/mastery and modesty/harmony presented in Section 6.6).

6.3 Baseline Simulation Model

This section introduces the baseline simulation model that is used throughout this chapter. This baseline simulation model is then explored and altered in the following sections for illustrating the various aspects of our theory. In order to get an overall idea of the content of this model, an overview of the properties of the various entities and their influence on each other is introduced in Figure 6.2.

This section introduces concepts and objects that belong to the environment in Section 6.3.1, and the coordination structure in Section 6.3.2. Then, the simulation loop is presented in Section 6.3.3. Finally, parts of agent decision processes that are shared by later implementations are presented in Section 6.3.4.

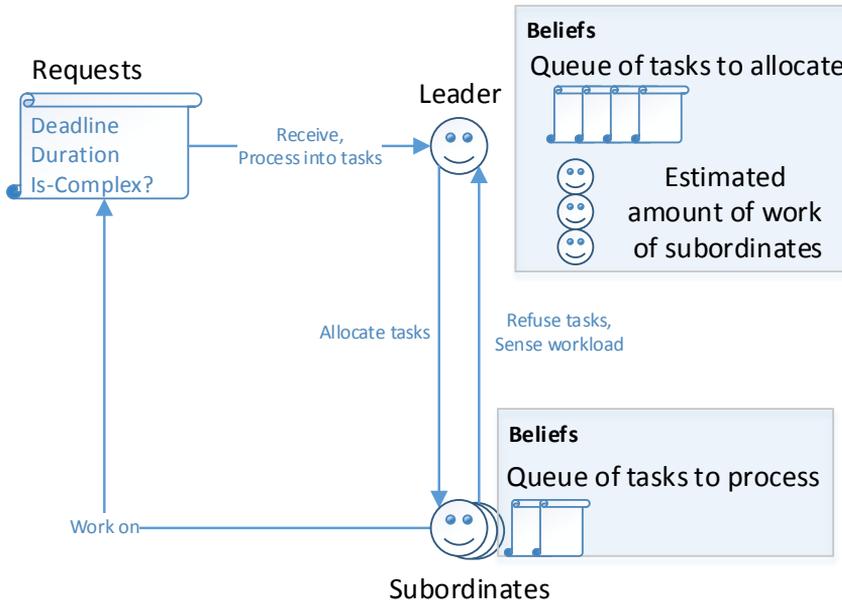


Figure 6.2 – Entities of DO-IT-2.0 and their relationships

6.3.1 Concepts & Objects

This section introduces core concepts and objects of the simulation model. The evolution of these objects over time is represented by the simulation loop in Section 6.3.3.

Environment

The environment contains the variable *time*, which measures the amount of time that goes on. *time* counts the number of cycles and corresponds to one hour of work-time.

The environment also stores the set T of tasks to be processed. Each task possesses a deadline, represented by $deadline : T \rightarrow \mathbb{N}$. The amount of work to be performed for completing a task is represented by $remaining : T \rightarrow \mathbb{N}$. The value $remaining(t)$ decreases when agents work on t .

When a task is not handled in due time, it can become unfeasible and is failed by the organization. Formally, a task t is *unfeasible* if $time(t) + remaining(t) > deadline(t)$.

Tasks can be complex and thus require extra unexpected processing time. The complexity of a task is modeled by: $is_complex? : T \rightarrow \{\top, \perp\}$ where $is_complex?(t) = \top$ means that t is complex. Otherwise, $is_complex?(t) = \perp$ means that t is simple. Agents cannot perceive the complexity of a task.

A task is *completed* if $remaining(t) = 0$ and $is_complex?(t) = \perp$. If $remaining(t) = 0$ but $is_complex?(t) = \top$, then the task gets expanded, as presented in Section 6.3.3.

Task Generation

In order to model the various environments presented in Section 6.2, we need to be capable of determining various patterns for generating tasks. The generation mechanism is characterized by the following variables.

First, we want to be capable of manipulating the number of incoming tasks at a given time. This manipulation is represented by $\#tasks : \mathbb{N} \rightarrow \mathbb{N}$, where $\#tasks(k)$ is the number of tasks to generate in round k . Second, the duration of tasks can vary, represented by the variable $base_duration \in \mathbb{N}$. Third, the time at which the deadline is set can also vary, determining the amount of time remaining for completing the task. This deadline is represented by $deadline : T \rightarrow \mathbb{N}$. Environmental variability corresponds to the (visible) variation with regard to the standard duration. Formally, environmental variability is represented by $variation \in \mathbb{N}$. Tasks have an initial duration randomly drawn from $[base_duration - variation; base_duration + variation]$. Finally, environmental complexity corresponds to the relative risk that incoming tasks are more complex than expected. Formally, this complexity is represented by $environmental_complexity \in [0, 1]$. New tasks have the probability $environmental_complexity$ that $is_complex?(t) = \top$. The rest of this section further describes how these environments are concretely implemented and used for running simulations.

Intensive Environment: In this environment, tasks come at a high pace and have to be processed quickly. Inefficiency in allocating tasks or task-less subordinates lead to failures. $\#tasks(n)$ is set to 1, for all $n \in \mathbb{N}$. $base_duration$ is set to 10. The deadline is set to $time + remaining + 2$. $environmental_complexity$ is set to 0 and $variation$ is an experimental variable. The simulation is run for 500 rounds, this duration is the same for all simulations in this chapter.

In a real-world scenario, this environment is the environment of, for example, a service in charge with setting up computers. In order to avoid stocking computers (saving space and risks for burglary), the organization works on “just-in-time” delivery. This task is relatively simple even if variations can occur, due to the different computational power of different devices or the amount of setup to perform.

Complex Environment: This environment corresponds to a basic environment where environmental complexity can evolve. In this environment, tasks are steadily generated and are sufficiently easy to solve, thus relatively inefficient allocations can occur without immediately leading to failures. Formally, $\#tasks(n)$ is set to 1 if n is even and 0 otherwise. $base_duration$ is set to 15. The amount of time remaining for performing the task t , $deadline(t)$ is set to 30. $variation$ is set to 0 and $environmental_complexity$ is an experimental variable.

In a real-world scenario, this environment is the environment of a typical helpdesk. With such a helpdesk, the organization receives tasks relatively regularly, without limited propensity for bursts or tight deadlines. Nevertheless, the range of incoming tasks is wide, potentially leading to complex resolutions (e.g. chains of problems).

Irregular Environment: In this environment, tasks come in busts or jerks, many at once. The organization has enough work-force for coping with all incoming tasks, but leaders have limited room for making mistakes. In this environment, every 20 time units a batch of $k \in \mathbb{N}$ tasks is generated at once, where k is an experimental variable. $remaining(t) = 5$ and $deadline(t) = 20$, so these tasks are relatively easy to process.

variation and *environmental_complexity* are set to 0.

As a real-world scenario, this environment can be seen as, for example, the work of the customer service of a company regularly releasing software updates. For each new update, the organization has to help many clients at once to update successfully their system.

6.3.2 Coordination Model

The coordination mechanism is based on an organization, more precisely on a simple structure. This organization has two roles: one leader and (several) subordinates. Leaders are in charge of allocating tasks to subordinates and subordinates have to perform allocated tasks. Figure 6.3 illustrates the interaction loop.

The set of agents is represented by A . The leader is represented by $a_l \in A$. Subordinates are represented by the set $S = A \setminus \{a_l\}$. In our simulations, $|S|$ is set to 10.

The coordination structure provides a queue of tasks to agents. This queue details the tasks to process and the order in which they have to be processed. Formally, this queue is represented by $queue_a$ of tasks in T for each agent $a \in A$. The next task to be processed is the first task of $queue_a$. This task is represented by $next_task_a$.

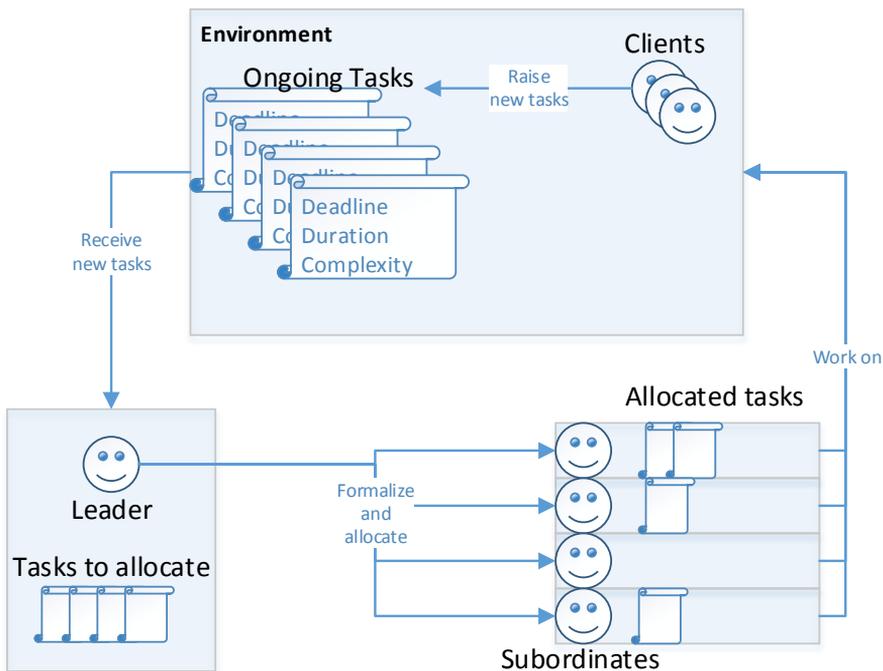


Figure 6.3 – Overall processing loop of DO-IT-2.0

6.3.3 Simulation Loop

The simulation loop formalizes the evolutions of agents, organizations and environments. As for DO-IT, each round represents one hour of work. Algorithm 2 introduces the sequence of steps to be executed by our simulation.

Following subsections further detail how the model evolves during each step.

Initialization

The environment contains no tasks initially. Agent queues are empty.

Generating New Tasks

Each round, $\#tasks(time)$ tasks are generated. $duration(t)$ is randomly drawn from $[base_duration - variation, base_duration + variation]$ and $deadline(t)$ is set similarly. $is_complex?(t)$ is set to \top with a probability of $environmental_complexity$.

New tasks are added at the end of the leader's queue $to_allocate_{a_l}$.

Algorithm 2 Simulation Update Cycle

```

1: Initialize agents, the organization and the environment
2: for the desired number of iterations do
3:   Increment time
4:   Generate new tasks and give them to the leader
5:   Fail and remove unfeasible tasks
6:    $d_l \leftarrow \text{decision}(a_l, \text{how to allocate tasks?})$ 
7:   if  $d_l = \text{evaluate \& allocate}$  then
8:      $a_l$  monitors a subordinate
9:     if  $queue_{a_l}$  is not empty then
10:       $a_l$  evaluates the duration  $next\_task_a$ 
11:       $allocate\_task(next\_task_a)$ 
12:     end if
13:   else
14:     for two times do
15:       if  $queue_{a_l}$  is not empty then
16:          $allocate\_task(next\_task_a)$ 
17:       end if
18:     end for
19:   end if
20:   for all  $a_s \in S$  do
21:     Work on first task
22:   end for
23: end for

```

Algorithm 3 *allocate_task(t)*

```

1:  $a_s \leftarrow \text{decision}(a_l, \text{who to allocate } t \text{ to?})$ 
2: if  $a_s = \perp$  then
3:   exit this procedure
4: else
5:   if  $\text{decision}(a_s, \text{accept } t?) = \top$  then
6:     allocate  $t$  to  $a_s$ 
7:   end if
8: end if

```

Remove Unfeasible Tasks

Tasks that can obviously not be accomplished are removed from agent “to_process” queues. Leaders only remove only unfeasible tasks (when the deadline is closer than the time to completion). Subordinates remove the tasks that, after having performed tasks that are prior in the queue, will be unfeasible at the time the agent will handle them. Note that in this process, durations ignore whether tasks are complex or not.

Decision: How to Allocate Tasks?

The leader has to decide how to allocate tasks. Possible answers for this decision are: “evaluate & allocate” or “quickly allocate”. With the first option, the leader can monitor one subordinate, sense the amount of work remaining in the queue $next_task_{a_l}$ and perform one allocation. With the second option, leaders skip evaluation and have two opportunities for allocating tasks to subordinates, allowing potentially to allocate two tasks, each for a different subordinate.

The way agents make this decisions and following ones is detailed in Section 6.3.4 which describes agent decision processes.

Monitoring a Subordinate

The leader selects a subordinate and can monitor it. During its next decision, the leader will perceive the amount of work of this subordinate. Formally the amount of work is represented by:

$$\sum_{t \in queue_{a_s}} remaining(t)$$

For the sake of simplicity, the decision about which subordinate to monitor is represented by a round-robin process.

Evaluating the Duration of the Next Task to Allocate

The leader will obtain as a percept the expected duration of the task it is currently trying to allocate. Note that the leader ignores whether the task is complex.

Decision: Whether and Who to Allocate the Next Task?

The leader has to decide whether to allocate $next_task_{a_l}$ and it does, which subordinate to allocate the task to. The set of allowed actions for this decision is $\{allocate(a) | a \in S\} \cup \perp$, where \perp means “no allocation”.

Decision: Accept Proposed Tasks?

Subordinates that are proposed tasks have to decide whether they accept them or not. Formally, agents have two possible actions for this decision: $\{\top, \perp\}$, where \top means accept and \perp means reject.

Allocating tasks

The task moves from the top of leader’s queue to the bottom of the allocated subordinate’s queue.

Working on tasks

Each round, subordinates work on the tasks they are allocated. Formally, each subordinate $a_s \in S$ works on $next_task_{a_s}$, if $queue_{a_s}$ is not empty. In this case, $remaining(next_task_{a_s})$ is decreased by 1.

If $remaining(t) \neq 0$, nothing happens. Else, if $remaining(t) = 0$ and $is_complex(t) = \perp$, then the task is completed and removed from the queue. Else, $remaining(t) = 0$ and $is_complex(t) = \top$, then the task is found to be more complex than expected and thus require more work. In that situation, $remaining(t)$ is set to 30 (more work), $deadline(t)$ is increased by 40 (extended deadline) and $is_complex?(t)$ is set to \perp (the task is no longer complex).

6.3.4 Agent Decision Process

This section introduces how agents make decisions. This agent decision process is compliant with our decision architecture from Chapter 4.

Our decision model aims at being as simple as possible because, this model is to be used for making social simulations. As advocated in Epstein (1999), these simulations are aimed at explaining how individual decisions can lead to collective outcomes. In order to produce coherent explanations, avoiding noise in results and keeping the model computationally viable, individual decisions should remain as simple as possible.

To that extent, in this chapter, we implement the agent architecture introduced in Chapter 4, decision rules as simple as possible.

Memory Structure

This structure introduces the static elements of agent memory that remain present throughout the decision cycle.

Values: We consider two values: *autonomy* and *obedience*. We consider two value systems: *obedience* > *autonomy* and *autonomy* > *obedience*. The influence of these two systems is directly integrated within agent decisions as decision rules. For instance, we can say “if obedience is more important than autonomy, the agent prefers the option. . .”.

Beliefs: Agents store accurate information about *time*.

They can also store beliefs about tasks. For a given task $t \in T$, each agent $a \in A$ can store beliefs about the amount of remaining work and its deadline. This information can be used for determining whether a task is unfeasible, whether a queue of tasks is unfeasible, what is the first tasks to fail in an unfeasible queue and whether adding a new task leads to an unfeasible situation.

Leaders can store beliefs (an integer) about the estimated amount of remaining work to be performed by each subordinate.

Desires & Intentions: While a complete formalization of these aspects can be proposed for matching the model from Chapter 4, we prefer to leave it out here. Detailing this additional machinery is not useful with regard to the properties we investigate and the relative simplicity of the environment. Nevertheless, these decision elements are much more relevant in more complex environments, where agents have to keep up with objectives or interact with others about detailed environmental information.

Decisions

This section introduces concrete decision rules for agents in the baseline model. For sake of simplicity, we assume that agents keep their beliefs up to date during the percept processing phase: subordinates receive (accurate) information about tasks they have in their queue or proposed tasks. Leaders receive (accurate) information only when actively evaluating a task or a subordinate.

For each decision, we introduce first how rationality can lead to a decision or remove irrelevant choices. If rationality is insufficient for making the decision, we introduce available decisions that raise a dilemma since none strictly dominates the other. Then, we introduce how culture supports the resolution of this dilemma.

Decision: How to Allocate Tasks? Leaders have to decide which strategy they pursue for allocating tasks. Two solutions are available: (1) quickly allocating, granting them with two attempts and (2) spending time for checking the situation before allocating once.

Dilemma: Both choices are rationally acceptable. In comparing the two, decision (1) supports more efficient allocation but grants less information and thus less precision than decision (2).

Influence of Culture: These decisions can be directly related to the values of autonomy and obedience. Autonomy supports option (1): this option is more efficient than the latter and does not imply watching subordinates. Obedience/social order supports option (2): the leader is assumed to be in charge of the group and thus should gather better information for making appropriate decisions.

Decision: Who to Allocate? Leaders have to determine whether they allocate the next task and who to allocate it.

Rational Decision: The target of the allocation is an agent with the least amount of work according to leader's beliefs. If, due to quick allocations, the leader does not keep track of the amount of work performed by subordinates, then the target is the subordinate

that has not been allocated a task for the longest amount of time. Formally, this decision is based on a round robin.

Dilemma: A dilemma occurs when the leader knows that the subordinate is busy but can still accomplish the task. In that case, by proposing immediately the task, the leader can make an early allocation. This early allocation can lead to unnecessary failures: if a former task of the subordinate may be more complex than expected, then the new task is likely failed. This failure could have been avoided by allocating on the last minute, when the subordinate is available for sure. However, waiting can also lead to inefficiencies, possibly triggering failures: maybe several subordinates will be available later at the same time, forcing the leader to have them wasting working time waiting for tasks.

Cultural Influence: In our simulations, we assume that the leader prefers to play safe by allocating tasks at the last minute. Nevertheless, this reaction is sensitive to other aspects of culture. As illustrated in Section 6.6, this decision change depending on values of assertiveness and modesty of the leader. In particular, leaders can prefer take risks and allocate tasks early, making sure that subordinates are always busy and thus avoiding inefficiency.

Accepting Tasks? Subordinates, when allocated tasks, have to decide whether they accept them or not.

Rational Decisions: The answer is always \top if the agent has no allocated tasks. The answer is always \perp if the agent is overloaded and this task can lead to failures.

Dilemma: A dilemma is raised when agents have already at least one task to perform. On the one hand, accepting another task may lead to failures if one of the already accepted task takes longer than expected. On the other hand, refusing the task incurs additional costs to the leader which can also lead to failures.

Cultural Influence: From the perspective of obedience, refusing orders is clearly unacceptable. From the perspective of autonomy, accepting and refusing are two valid options. In this baseline model, agents prefer to refuse in order to avoid failures. Additional values can better support this decision, as illustrated in Section 6.6 with values of assertiveness and modesty.

Note that if a subordinate a_s accepts the task $t \in T$ proposed by the leader and the leader has the belief that the work remaining for processing t is k , then the leader adds k to the estimated amount of work to be performed by a_s . Each round, the estimated of work for all subordinates is decreased by 1, with a minimum of 0.

Conclusion

This section introduces the baseline model that we use as focal point for the rest of this chapter. Following sections either study this model or alter it with different coordination mechanisms or cultural models, in order to illustrate concretely the various properties of our theory.

This model and its subsequent extensions are implemented by a NetLogo program (version 5.0.4). This programming language and development software is particularly practical for designing and running social simulations. As an illustration, Figure 6.4 presents a screenshot of the user interface.

In the light of the content of this model, we can better explain how it fits with regard to the goal of this chapter. This model supports rich interactions between leaders and subor-

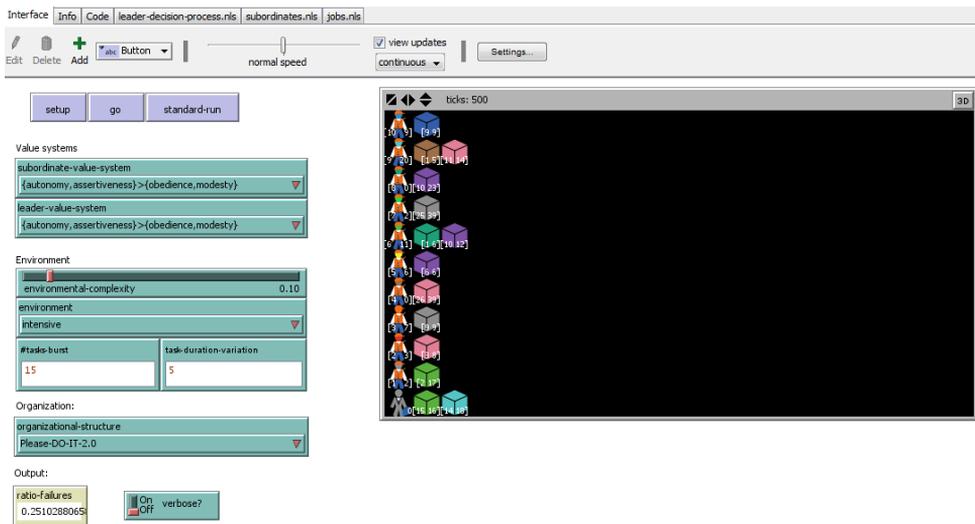


Figure 6.4 – Screenshot of the implementation of DO-IT-2.0 simulation in NetLogo 5.0.4

dinates while leaving out technical (individual-oriented) details. This simplicity allows us to keep a simple agent decision model while still maintaining several important dilemmas for which culture provides an answer.

6.4 Simulation 1: Exploring the Baseline Model

This section explores the behavior of the baseline model. This section aims at showing that this model is realistic and illustrates relationships introduced by our theory. Then, outcomes of this baseline are used as a point of comparison for investigating relationships proposed by our theory, in later sections.

This section is structured as follows. First, Section 6.4.1 aims at providing a basic validation of this model, by testing the dynamics of our model in different cultural conditions against core properties recognized by theories of culture and empirical findings. Then, Section 6.4.2 shows how this model illustrates our theory.

6.4.1 Experiments

In order to test whether our model makes sense with regard to reality, we compare it with some core findings from social sciences.

Our cultural model relies on values of autonomy and obedience. These two values are related to the PDI-CD as identified in Section 3.2. PDI-CD is related to several properties from which we investigate two most important ones in the following subsections.

Culture and Environmental Complexity

Theories relate PDI-CD with performance depending on environmental complexity. Organizations for which members share a high PDI-CD culture tend to achieve better performance in simpler environments. Conversely, organizations with low PDI-CD culture tend to handle better more complex environments. These properties are introduced in Hofstede et al. (2010a) and depicted in Figure 2.3.

In order to evaluate these two properties, we compare collective outcomes depending on the environment and underlying culture.

High PDI-CD Leads to Better Performance in Simple Environments: In order to investigate this property, we consider the “intensive” environment, presented in Section 6.3.1. This environment is simple, leaders can make adequate allocations without being surprised by tasks that are longer than expected. However, this environment is also unforgiving: mis-allocations generally lead to failures. In order to illustrate organizational adaptiveness, we investigate what happens when the duration of tasks varies.

Results are presented in Figure 6.5. These results indicate that for any degree of variation, organizations with a culture that favors obedience over autonomy (those with higher PDI-CD) achieve higher performance in this simple environment. These results match empirical findings from social sciences.

Low PDI-CD Leads to Better Performance in Complex Environments: In order to investigate increasingly complex environments that are not too overwhelming, we consider the “standard” environment and alter environmental complexity (probability that tasks are complex).

Results are presented in Figure 6.6. These results indicate that, when environmental complexity increases, organizations with a culture where autonomy is preferred over obedience (with lower PDI-CD) achieve higher performance. These results match empirical findings from social sciences as presented in Hofstede et al. (2010a), briefly introduced in Figure 2.3.

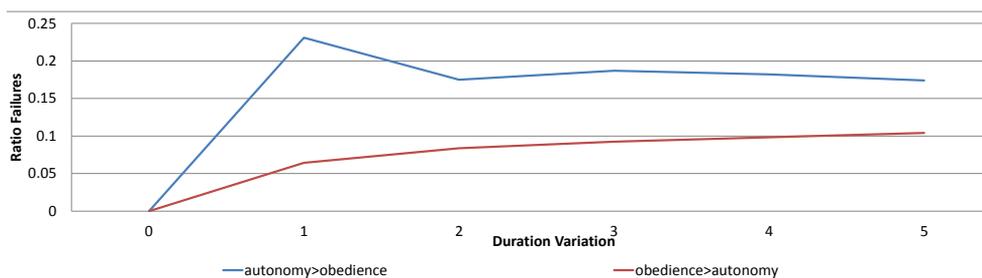


Figure 6.5 – Ratio of failed tasks (Y-axis) depending on the variability of the duration of incoming tasks (X-axis) in the intensive environment. Each curve represents a cultural preference towards either autonomy or obedience. These curves measure organizational efficiency and flexibility in simple environments depending on culture

Culture and Robustness

Hofstede et al. (2010a) tightly correlates PDI-CD with a tendency towards centralizing decision and information around leaders. One of the properties of centralization is a relatively low robustness. This section investigates whether our simulated organization achieves lower congestion robustness (lower performance when many tasks arise at once) depending on PDI-CD³. In order to verify this property, we investigate organizational performance in the irregular environment, presented in Section 6.3.1.

Results are presented in Figure 6.7. These results indicate that, when more tasks have to be processed at once, organizations with a culture that prefers autonomy over obedience (in lower PDI-CD) achieve higher performance and thus higher robustness to congestion. These results match empirical findings from social sciences.

Conclusion

Experiments presented in this section show that the output of our model matches core empirical findings from studies of cultures from the perspective of social sciences.

6.4.2 Back to Our Theory

This baseline model provides a concrete illustration of several relationships drawn by our theory. These links are highlighted further in this section.

Abstract, Coherent and Recurrent Interaction Patterns Supported by Culture

This model illustrates the relationship between different cultural settings and interaction patterns. Each cultural setting supports a specific *abstract* interaction pattern.

If autonomy is important, this culture supports an abstract interaction based on decentralization. Subordinates are in charge of managing themselves. This assumption is

³In this simulation, agents cannot fail. Thus, we cannot study the influence of culture on structural robustness (robustness to failing agents)

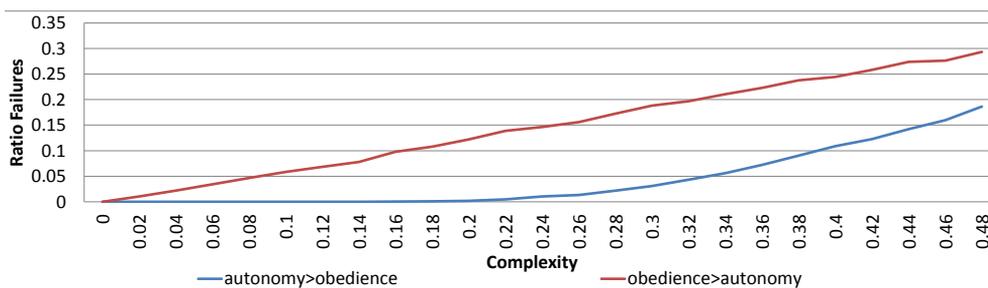


Figure 6.6 – Ratio of failed tasks (Y-axis) depending on the complexity of the input (X-axis) in the complex environment. Each curve represents a cultural preference towards either autonomy or obedience. These curves indicate organizational performance in complex environments depending on culture

taken into consideration when leaders delegate tasks to subordinates. Likewise, leaders avoid watching subordinates for not violating their privacy.

If obedience is important, then culture supports an abstract interaction based on centralization. The decision power is left to leaders, and subordinates are expected to comply with it. In our concrete organization, leaders are in charge of managing the workload of subordinates.

These interactions are *coherent*. If autonomy is important, both leaders and subordinates either directly expect and act towards or indirectly support that subordinates manage their own workload. If obedience is important, both leaders and subordinates directly expect or indirectly support that leaders manage the workload of subordinates. In both cases, managing the workload is done by exactly one entity (no redundancies and at least someone cares about it). Section 6.5 further illustrates the cultural origin of this coherence by showing incoherences resulting from cultural mismatches.

Recurrence is illustrated by the occurrence of similar interaction patterns in various environments. Furthermore, this interaction pattern is relatively abstract and implementable in a wide range of situations. Section 6.8 further illustrates recurrence in showing that similar interaction patterns occur in spite of changes in the coordination mechanism.

Culture, Coordination and Environment

This model illustrates the relationship between culture and coordination mechanisms.

This model shows that coordination mechanisms can support expression of values. Indeed, coordination mechanisms, by making explicit the roles of leaders and subordinates, provide a crucial basis for integrating the influence of autonomy and obedience on decisions. The other way around, in an organization without leaders and subordinates, autonomy and obedience are harder to relate to decisions that have to be made.

In a more general perspective, this coordination setting raises decisions that can be easily supported by autonomy and obedience. Furthermore, for these decisions, values provide sharp evaluations (e.g. “very good” or “very bad”), allowing to discriminate the various alternatives or cut early decision paths that are not value-compliant. For instance, complying with or refusing orders ranges from “very good” to “very bad” for obedience but has hardly an impact for the value of “modesty”.

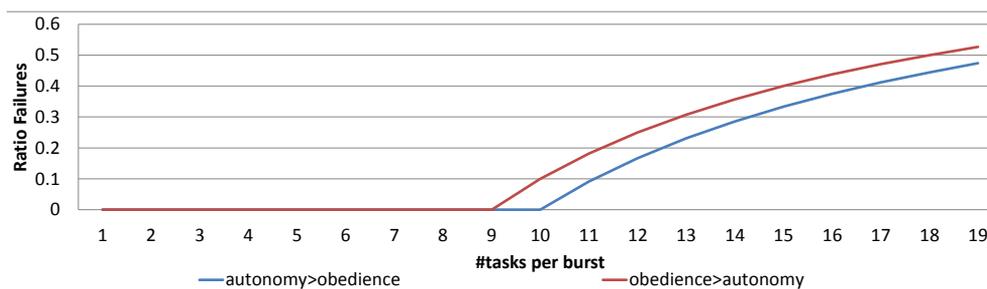


Figure 6.7 – Ratio of failed tasks (Y-axis) depending on the amount of tasks in each burst (X-axis) in the irregular environment. Each curve represents a cultural preference towards either autonomy or obedience. These curves indicate organizational robustness, depending on culture

This model also illustrates the influence of restrictions imposed by coordination mechanism on cultural expressiveness. In the current setting, the coordination mechanism restricts agent interactions to “allocations-acceptance”. This form of coordination prevents, for instance, some type of interactions and decisions supported by autonomy, such as direct interactions between subordinates and initiative-taking.

The relationship between coordination and cultural expressiveness is studied further in Section 6.7 by restricting agent freedom and in Section 6.8 by investigating a more open coordination structure.

The Impact of Cultural Richness on the Influence of Culture

This aspect is not illustrated by this baseline model, which considers similar cultures. Nevertheless, this baseline model offers a point of comparison for later experiments. On this track, Section 6.5 illustrates what happens when agents have less cultural background in common. Conversely, Section 6.6 illustrates what happens when agents share a richer cultural background.

These later models allow illustrating the properties of the impact of cultural richness on the influence of culture (e.g. more precise interaction patterns, consequences on collective outcomes).

From Individual Decisions to Interactions to Collective Outcomes

This baseline model is complete enough for illustrating one of the core purposes of our theory: explaining how the influence of culture on individual decisions leads to expected influence of culture on collective outcomes. As depicted in Section 6.1, this link can be explained in two steps: in relating individual decisions to interactions and interactions to collective outcomes.

From individual decisions to interactions: Autonomy supports decisions leading to decentralized interaction patterns, while obedience supports decisions centralizing power around leaders. This link has already been detailed in previous sections. If autonomy is important, leaders propose tasks to subordinates and expect them to decide by themselves to accept them or not. If obedience is important, leaders decide for subordinates the tasks they should accomplish and subordinates expect this decision to be (accurately) made by the leader.

From interactions to collective outcomes: Centralized systems allows a tighter management of subordinates than decentralized systems. Nevertheless, this management is limited by the capability of the leader for managing the group. This capability is limited by environmental complexity (need more time for making decisions) and predictability (outdated or inaccurate information). To that extent, organizations supported by a culture of obedience tend to achieve better efficiency and flexibility in simpler environments (easy to centrally handle) than organizations supported by a culture of autonomy. Nevertheless, this tendency is reversed in more complex environments (in our case, because information can be outdated). For similar reasons, centralization is more sensitive to (congestion) robustness (overloading the central agent).

Conclusion

This section studies the baseline model presented in Section 6.4 by checking that it makes sense compared with empirical findings from social sciences and by relating it to our theory. This example illustrates most of the aspects of our theory: culture supports an abstract, coherent interaction pattern. It also illustrates that the expressiveness of a culture is sensitive to coordination mechanisms and the environment. More importantly with regard to the goal of this chapter, it illustrates that the influence of culture on collective outcomes can be operationally explained in terms of the influence of culture on individual decisions and interactions.

The baseline model is then used in the following sections for as a point of comparison for highlighting the various relations proposed by our theory.

6.5 Simulation 2: Mismatching Cultures— Cultures, Coherence of Interactions and Collective Outcomes

This section illustrates two relations presented in our theory: (1) sharing a culture supports coherent interactions; (2) coherent interactions tend to supports better collective outcomes, but “accidental good performance” can result from incoherent interactions resulting from cultural mismatches.

In order to do so, we study interactions and collective outcomes in a situation where agents have completely mismatching cultural backgrounds. This situation is put in comparison with the baseline situation from Section 6.4 where agents all share the same one-value culture. Since in this baseline culture supports a collective preference towards autonomy or obedience, in this section we focus on the cases where interacting agents have different value-orders. In the case of the values of autonomy and obedience, these values directly influence the interaction between leaders and subordinates. To that extent, we focus on the case where leaders and subordinates have different value-orders.

This section is structured as follows. Section 6.5.1 relates cultural mismatches observed within our simulation with empirical findings from social sciences. Then, Section 6.5.2 relates these simulations with properties studied in our theory from Section 6.1.

6.5.1 Experiments

In order to check the match between this model and reality, we investigate whether this model reproduces findings of theories of culture.

Theories of culture dedicate a lot of attention to the description of the influence of cultural mismatches on interactions (e.g. an extensive text in Hofstede et al. (2010a) is dedicated to cross-cultural encounters). They describe that cultural mismatches can raise “cultural clashes”, or misunderstandings, mostly through examples. Nevertheless, they do not provide encompassing explanations about how these clashes occur. They suggest that these clashes damage collective outcomes in situations of coordination but they do not detail to which extent (this aspect is detailed by our theory and illustrated through experiments in Section 6.5.2). In this section, we relate the most relevant example proposed by these theories with the output of our simulation model.

In this example from Hofstede et al. (2010a), a high PDI-CD leader (promoting obedience) who gives an order to a low PDI-CD subordinate (promoting autonomy). The leader expects the subordinate to accept the order immediately, but the subordinate asks instead for more information about the reason of this order. Our model replicates similar situations. A leader promoting obedience can allocate a task to a subordinate promoting autonomy. If this subordinate is already busy, the subordinate can refuse the task. This situation goes against expectations of both agents. The decision of the leader is based on the expectation that the subordinate must accept (this is why the leader makes careful allocations). The decision of the subordinate is based on the expectation that instructions of the leader are “proposals” and not “orders”.

6.5.2 Back to Our Theory

This simulation model illustrates two relations proposed by our theory: sharing culture supports coherence and coherence provided by culture generally but not necessarily supports good collective outcomes.

Culture and Coherence

Our theory suggests that sharing a cultural background supports coherent abstract interaction patterns. In Section 6.4.2, we show that sharing preferences towards values of autonomy or obedience supports coherent interaction patterns (e.g. centralization and decentralization). In comparison, this section shows that when agents share less cultural background (here, values), then their interactions become less coherent. In order to do so, we investigate situations where leaders and subordinates have different cultural values.

If leaders prefer obedience and subordinates autonomy, then both leaders and subordinates assume that they are in charge of managing the workload of subordinates. This mismatch entails redundancies in terms of workload management. These redundancies can lead to disagreements (e.g. because both expect to be empowered, as illustrated in Section 6.5.1) and can also lead to inefficiencies (slow allocations and possible rejections).

If leaders prefer autonomy and subordinates prefer obedience, then both expect the other to manage the workload of subordinates. As a consequence, workload is completely unmanaged. This lack of management can easily raise dramatic failures (locally overloaded of subordinates, while others are idling).

In both cases, interactions resulting from mismatching culture are incoherent. Agents are culturally driven towards different interaction patterns than those that are occurring. Thus, agents create erroneous expectations about behaviors of others. From an external perspective, such an interaction pattern can be hardly considered as supportive for coordination.

Culture and Collective Performance

Our theory posits that when individuals in a group share the same culture, they tend to achieve relatively good performance in general. Nevertheless, we do not claim that sharing culture necessarily improve performance in any situation. First, because specific cultures may be particularly inappropriate for handling some coordination problems. Second because cultural mismatches can sometimes lead “by accident” to better outcomes

than when culture is shared. Figure 6.8, Figure 6.9 and Figure 6.10 introduce plots expanding the experiments conducted with the baseline model in adding cases where leader’s and subordinates’ culture mismatch.

These plots indicate sharing the same culture generally leads to good outcomes. The mismatching culture, where leaders prefer autonomy and subordinates prefer obedience, generally leads to poor performance in comparison with other cultures (apart from the irregular environment where good performance is reached “by accident”, further explained in the following paragraphs). The mismatching culture where leaders prefer obedience and subordinates prefer autonomy, displays surprising results, because they are comparable with those of shared culture. These results, seemingly going against our theory, are actually caused by simplifications in our model. Indeed, evaluating the duration of tasks is free while it should take time. Plus, disagreements between leaders and subordinates do not raise conflicts while they should. If evaluating the duration tasks have had a cost, then redundant evaluations raised by this mismatching culture would have lead to dramatic failures in the intensive environment where time cannot be wasted. If conflicts can occur, disagreement about who should manage subordinate’s workload would also lead to conflicts and failures in complex environments. Thus, this simplification makes that the results of these mismatching cultures are “optimistic”. Their performance should be worse than matching culture.

Figure 6.10 highlights that cultural mismatches can lead “by accident” to higher performance. This figure illustrates that when leaders favor autonomy and subordinates favor obedience, then higher performance is obtained in the irregular environment. Thus, this cultural mismatch leads, “by accident” to higher robustness.

Further explanations can be obtained in studying interactions in more details. In the irregular environment, failures occur when inefficient allocations occur, due to slow allocations or rejections from subordinates. In the case of our mismatching culture, these inefficiencies are avoided by the blind allocation and acceptance from both sides. Subordinates, by never refusing orders from leaders, take a high risk of failure but do not cost

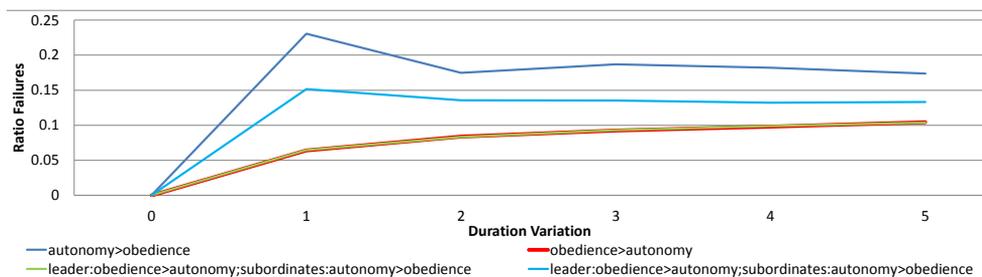


Figure 6.8 – Ratio of failed tasks (Y-axis) depending on the variability of the duration of incoming tasks (X-axis) in the intensive environment. Each curve represents an (un)shared cultural preference towards autonomy or obedience. When unshared, the preference order of leaders towards autonomy and obedience differs from the one of the subordinates. These curves indicate the organizational efficiency and flexibility in simple environments depending on culture.

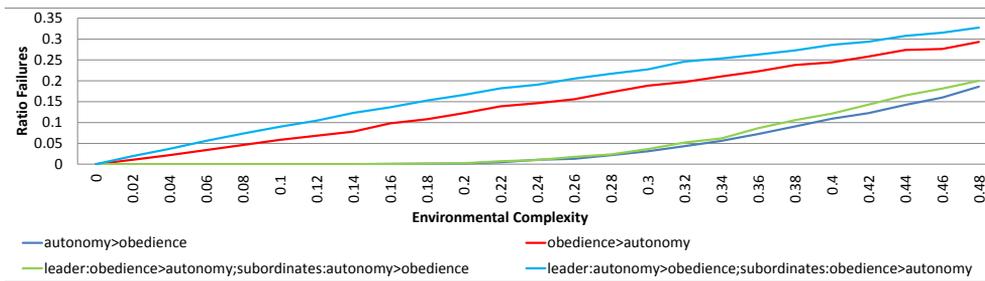


Figure 6.9 – Ratio of failed tasks (Y-axis) depending on the risk that incoming tasks are complex (X-axis) in the complex environment. Each curve represents either an (un)shared cultural preference towards autonomy or obedience. These curves indicate the organizational performance in complex environments depending on culture

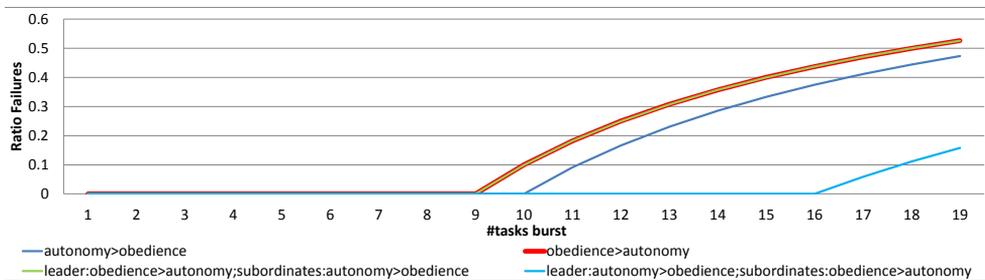


Figure 6.10 – Ratio of failed tasks (Y-axis) depending on the amount of tasks in each burst (X-axis) in the irregular environment. Each curve represents either an (un)shared cultural preference towards autonomy or obedience. These curves indicate the organizational robustness, depending on culture (note that green and red curves overlap because in both cases leaders carefully allocate tasks and have accurate information about who is available)

time refusing tasks. “By chance” this blindness works in this precise environment. Nevertheless, this blindness is hardly adequate for achieving high performance in a wide range of situations.

Conclusion

This section illustrates through a concrete example the relationship presented in the theory linking cultural mismatches, coherent interaction and collective performance. Culture appears to support coherent interaction patterns which in turn lead to relatively good performance in general. We also show that this support is not as omnipotent as suggested by theories: cultural mismatches and misunderstandings can sometimes lead to better performance.

6.6 Simulation 3: Richer Culture—Relating Cultural Richness, Richer Interactions and Collective Outcomes

This section illustrates the influence of the richness of cultural background on interactions and collective outcomes described by our theory. This section complements the illustration from Section 6.5 by showing the influence of richer shared cultural backgrounds instead of focusing on cultural mismatches (representing more sparse cultural backgrounds).

In order to do so, we expand the cultural model from the baseline in adding two additional values: assertiveness and modesty. Then, we explore the match this richer model and empirical findings before showing how it illustrates our theory.

6.6.1 Simulation Model

This simulation model expands the baseline model with two values: assertiveness/mastery and modesty/harmony.

Overview

Assertiveness/mastery and modesty/harmony are two values from the Schwartz cluster of values (Schwartz (2006b)). As presented in Section 3.2, these values are strongly related to the two extremes of the MAS-CD. These values influence decisions about failure and individual performance related to the management of uncertainty.

Assertiveness/Mastery value focuses on evaluating good individual performance. This preference towards high performance can lead, in particular, to risk-taking attitudes. In our organizations, agents giving importance to assertiveness are more likely to allocate or accept tasks as long as they believe that all tasks can be achieved in due time. This early allocation and acceptance allows avoiding subordinates to wait for tasks. However, early acceptance can also support the risk of larger failures if an accepted task is more complex than expected.

Modesty/Harmony value negatively evaluates the risk-taking and causing failures, which is a direct source of disharmony. This desire for avoiding failures leads to a preference towards making safe decisions and avoiding early commitment. In our organizations, agents giving importance to modesty/harmony tend to avoid allocating or accepting tasks unless they are certain that the task can be processed. Given the duration required for processing unexpectedly long tasks due to complexity, agents allocate or accept tasks only when they have currently no task.

Value System

Formally, as for obedience and autonomy, assertiveness and modesty are evaluated by “good” and “bad”. In considering value systems, we consider the four following pairs:

- $\{obedience, assertiveness\} > \{autonomy, modesty\}$
- $\{obedience, modesty\} > \{autonomy, assertiveness\}$

- $\{autonomy, assertiveness\} > \{obedience, modesty\}$
- $\{autonomy, modesty\} > \{obedience, assertiveness\}$

This set of value-systems contains all possible value-systems which order the two pairs of antagonist values: autonomy vs. obedience ; assertiveness vs. modesty. We do not order non-antagonist values with each other because these values influence non-contradicting aspects of decisions.

Agent Decision Processes

This section introduces how new values influence dilemmas raised by decision processes and thus affect agent decisions. The decision process introduced in the baseline actually corresponds to the case where “*modesty > assertiveness*”⁴. In this section, we introduce what happens if *assertiveness > modesty*.

How To Allocate Tasks?: Values of assertiveness and modesty have no visible impact on this decision. For both values, pros and cons of each possible decision are relatively balanced. For assertiveness/mastery, quick allocation is more efficient but less accurate. For modesty/harmony, quick allocation may raise failures due to lack of accuracy but spending more time may raise failures as well due to lack of speed in allocating. For this decision, assertiveness and modesty are taken over by autonomy and obedience which are more relevant for this decision.

Who To Allocate Tasks to?: Values of assertiveness and modesty impact on the resolution of the dilemma related whether to allocate a task or not. If *assertiveness > modesty*, leaders accept to take the risk of performing early allocations in order to avoid later failures due to inefficiencies. In this case, they still propose tasks subordinates with the lowest believed amount of work. However, this time, they accept to do so even if the subordinate is already processing a task as long as the subordinate can accomplish this task in due time given their estimated workload.

Accepting Tasks?: The influence of obedience takes over values of assertiveness and modesty. If obedience is important, subordinates comply with orders and always accept tasks. Else, if autonomy is important, subordinates decide by themselves, under the influence of assertiveness and modesty. If *assertiveness > modesty*, they accept tasks as long as they believe they can fulfil them given their current queue, expecting that already accepted tasks are not complex. If they prefer modesty, they avoid taking risks by accepting tasks only if they have no allocated tasks.

6.6.2 Experiments

Hofstede et al. (2010a) indicate that a culture promoting assertiveness (high MAS-CD) tends to handle better larger amounts of tasks than one promoting modesty (low MAS-CD). A culture promoting modesty tend to be more adequate for handling more complex tasks, requiring adaptation to uncertain environments.

⁴This property is particularly interesting to note and is crucial for understanding how culture can support coordination, further detailed in Section 6.6.3 and Section 7.1.2

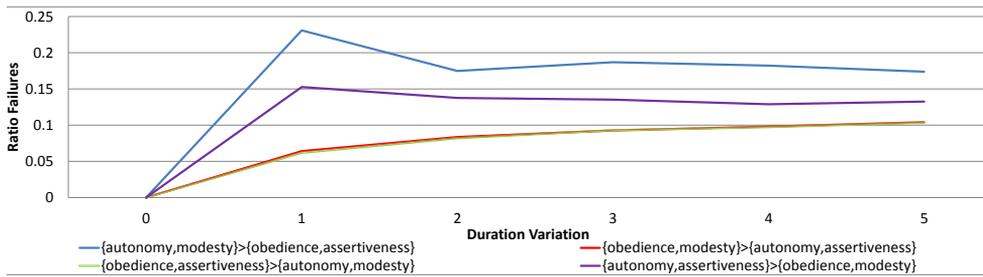


Figure 6.11 – Ratio of failed tasks (Y-axis) depending on the variability of the input (X-axis) in the intensive environment. Each curve represents a cultural preference towards autonomy vs. obedience and assertiveness vs. modesty. These curves indicate the organizational efficiency and flexibility in simple environments depending on culture.

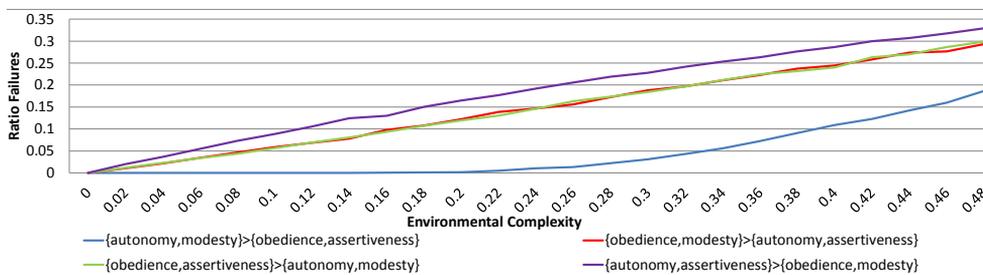


Figure 6.12 – Ratio of failed tasks (Y-axis) depending on the complexity of the input (X-axis) in the complex environment. Each curve represents a cultural preference towards autonomy vs. obedience and assertiveness vs. modesty. These curves indicate the organizational performance in complex environments depending on culture

As an experimental protocol, we investigate the performance of the organization through the intensive, complex and irregular environment in altering cultures with assertiveness and modesty.

Results are presented in Figure 6.11, Figure 6.12 and Figure 6.13. These figures indicate that a preference towards assertiveness leads to better performance when handling heavy workloads. In both intensive and irregular environments, for each cultural preference towards either autonomy or obedience, the culture for which assertiveness is more important than modesty achieves better results than when modesty is more important than assertiveness. Conversely, in complex environments, the culture promoting modesty achieves better performance. In more details, the culture for which modesty is more important than assertiveness achieves better results than their equivalents where this preference is reversed.

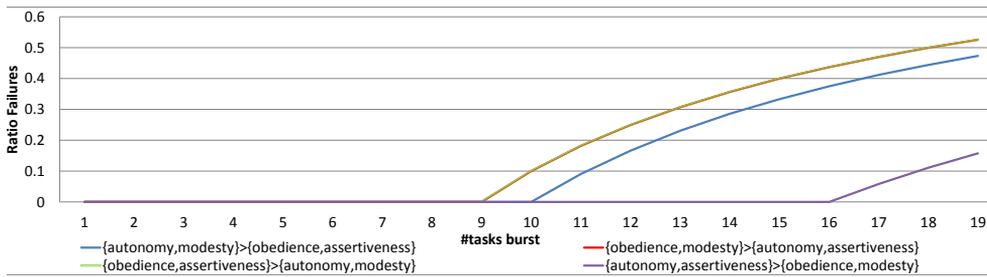


Figure 6.13 – Ratio of failed tasks (Y-axis) depending on the amount of tasks in each burst (X-axis) in the irregular environment. Each curve represents a cultural preference towards autonomy vs. obedience and assertiveness vs. modesty. These curves indicate the organizational robustness, depending on culture

6.6.3 Back to Theory

This model illustrates several relationships presented in our theory. In particular, it relates cultural richness to occurring interaction patterns and resulting collective outcomes. In addition, it supports the relationship between cultural richness and culture expressiveness.

Culture Supports Coherent Interaction Patterns

This example illustrates that each culture supports specific interaction patterns. Given the relative lack of correlation of the influence of each antagonistic couple of values on decisions, interaction patterns promoted by assertiveness and modesty can be considered relatively independently from the one promoted by autonomy and obedience⁵. Assertiveness promotes early commitment when making decisions. Conversely, modesty promotes last-minute decision, avoiding early commitment that one cannot hold. In combining these values, we can observe four types of interactions: “last minute allocation with centralized workload management”, “last minute allocation with decentralized workload management”, “early allocation with centralized workload management”, “early allocation with decentralized workload management”.

As before, when agents share the same culture then visible interaction patterns are internally coherent. Shared culture supports coherence between leader and subordinates with values of autonomy and obedience, as explained in Section 6.4.2. In addition, culture supports coherence in terms of balancing workload between subordinates. Assertiveness supports a collective agreement towards early allocation and modesty supports a collective agreement towards last-minute allocation. When these values are not culturally shared, incoherent interaction patterns can occur, particularly when autonomy is important: more assertive individuals accept the tasks that less assertive ones refuse. As a consequence, assertive individuals take more risks, taking tasks away from more modest individuals which waste more time waiting for tasks when available.

⁵This is not a general rule; some values can be tightly related. Resulting interaction patterns may be hard to express in terms of orthogonal influences

This model offers an opportunity for showing a crucial aspect of the support that culture can provide for coordination. In looking back to the baseline model, this model matches the case where, culturally, modesty is more important than assertiveness. During our effort of producing the baseline model, we obviously tried to avoid (implicit cultural?) conflicts when designing agents. We were looking for a coherent social paradigm. By default, we implicitly picked the “modesty > assertiveness” solution. Somehow, we thought this solution was the best, without considering in details the “assertiveness > modesty” solution (which is appropriate as well). In some sense, our own values implicitly driven us towards preferring the first solution against the latter.

This simple example hints to how culture can be used as a design tool for supporting coordination. What if we were working with other designers with different values? They could have as well selected the other direction, leading to coordination failures as depicted above in this section. Thus, in making culture explicit, system designers can highlight this distinction and avoid subsequent failures. In some sense, system designers can indicate the *social paradigm* (or underlying social atmosphere) of the whole system. This aspect, while difficult to define is crucial for achieving coordination and can be strongly captured by culture.

Cultural Richness and Interaction Patterns

This example illustrates the influence of cultural richness on interactions. In comparing with the baseline, the model of culture presented in this section is richer (agents share more cultural background). In turn, this richer cultural background drives more precise interactions. Indeed, this cultural background supports two variables of the leader-subordinate interaction: supports early vs. last-minute allocations in addition to leader-centered or subordinate-centered management of the workload. The wider is the shared cultural background, the more interactions are culturally-driven.

Cultural Richness and Culture Expressiveness

This example illustrates the influence of cultural richness on culture expressiveness. Values of assertiveness and modesty enrich the set of situations in which cultures influence agents. More precisely, these two values support decisions related to performance and uncertainty concerning failures. These decisions aspects add to the identity decision aspect influenced by values of autonomy and obedience. Thus, assertiveness and modesty can influence several decisions not supported by autonomy and obedience and vice versa (e.g. deciding how to allocate tasks).

Cultural Richness and Collective Outcomes

This example illustrates another property of culture. The influence of both value-systems appear to add their influence on collective outcomes. Furthermore, the influence of each value-system is independent on the environment. For instance combining autonomy and modesty supports the best collective outcomes in terms of robustness and capability for handling complex environment. These outcomes match theories of culture when taking each cultural dimension independently.

This capability of combining values is useful when aiming at using culture for supporting coordination. First, because the support provided by adding a given value can be strongly expected *a priori*, before having to implement these values. Second, because the support provided by multiple (matching) values appear to add to each other without seemingly introducing too many side effects.

Conclusion

This section illustrates the impact that richer cultural background has on resulting interactions and collective outcomes. In conclusion, the more culture is shared, the sharper expectations about other agents can be made. From the perspective of designing systems, using richer models of culture allows to drive more precisely agents, interactions and collective outcomes.

6.7 Simulation 4: Restrictive Coordination Mechanisms— Relating Culture and Coordination Mechanisms

The model presented in this section illustrates the influence of coordination mechanisms on cultural expressiveness. In order to do so, we expand the coordination mechanism presented in the baseline with more rules and show how it affects the expressiveness of culture.

6.7.1 Simulation Model

The simulation model is based on DO-IT-2.0, the only difference being to expand the rule base of the organization. The new organization is referred to as Rule-DO-IT-2.0.

The new rules of this organization force leaders to perform quick allocations, by using a round-robin. Subordinates are forced by rules to accept allocated tasks.

Considering agent decisions, the new coordination mechanism restricts agent decisions by leaving them with only one allowed action.

6.7.2 Experiments

Rules of these organization enforces the same behavior as the case *autonomy* > *obedience* for leaders, and *obedience* > *autonomy* for subordinates. Related outcomes for these decision rules are depicted by the adequate curve in Figure 6.8, Figure 6.9 and Figure 6.10. Nevertheless, since this decision rule is external to the agent, agent behavior does not change for all possible cultures. In other words, the performance of this organization is culturally-independent.

This cultural independence matches what can be expected in human societies. Culture can hardly play a role on interactions when the system is too regimented (e.g. forms and softwares in bureaucracies, pushing buttons in a production line).

6.7.3 Back to Theory

This example illustrates how coordination mechanisms can restrict the expression of culture. In the current model, agents have only one available choice. As a consequence, culture cannot be expressed.

A slightly more open rule system (e.g. leaving leaders free to decide how to allocate tasks) can support the expression of culture by including points in which agents have multiple choices to select from and selecting an option against another raises dilemmas.

As a side note, this lower expressiveness impacts on the range of possible collective outcomes that can be reached by the system. Indeed, constraints on agent decisions constrain possible interactions that constrain in turn possible collective outcomes.

Conclusion

This section illustrates the tight coordination relationship between the influence of culture on coordination mechanisms. Coordination mechanisms highlight situations of interactions and conceptually enrich decisions. Both raise conceptually-rich situations of interactions, the appropriate ground for expressing culture. However, coordination mechanisms can also be repressive (e.g. restricting too much freedom) and thus preventing the expression of culture.

This balance raises an important question to investigate when using culture for supporting coordination. When they want culture to be expressed, system designers should take these aspects into consideration when designing other coordination mechanisms.

6.8 Simulation 5: Changing Coordination— Relating Culture, Coordination Mechanisms, Recurrent Interactions and Collective Outcomes

This section illustrates how culture supports the occurrence of recurrent interaction patterns, which are relatively independent from external factors such as the environment and coordination. In order to do so, we consider a variant of the coordination structure from the baseline and show that similar interaction patterns result from similar cultural influences.

6.8.1 Model

This section introduces a new organizational structure, called DO-IT-Together-2.0, based on cooperation between agents.

Target Environment

Agents are allowed to work together in order to solve tasks. With some effort (spending 5 time units) agents can split a task in subtasks (in our model 5 subtasks). The remaining work of each subtask is only a portion of the initial duration (1/5). Tasks can be split only once. Each chunk can be complex, according to the standard *environmental_complexity*. A task has failed if any of its chunks is not finished before its deadline.

Coordination Model

The leader-subordinate structure is still available for supporting decision making. Nevertheless, in this organization, any agent can exchange tasks with others, including subordinates between them.

Decision Model

Leaders: If autonomy is important, leaders directly allocate tasks to subordinates, which in turn are in charge of dividing and sharing them.

If obedience is important, leaders assume that they are in charge of managing coordination of subordinates. To that extent, they take on themselves the charge of dividing the task and tightly allocating subtasks.

Subordinates: If autonomy is important, subordinates consider that they should manage their workload. They decide on their own to decide to split incoming tasks and to manage their own workload. From a modeling perspective, if their first task is too long and can be divided, they divide it. If they have several chunks of the same task to process, they allocate the first to another randomly drawn subordinate.

Subordinates giving high importance to obedience assume that leaders allocate them tasks they should be capable of performing. Furthermore, they do not want to mess with the information that leaders have by making local allocations. They try to perform tasks without dividing them. If the task is too long, they refuse it.

6.8.2 Experiments

Theories relate the influence of culture with collective outcomes relatively independently on a given coordination setting. Thus, our model should obtain the same cultural influences as those acknowledged in the baseline, presented in Section 6.4.1: (1) simpler environments should be better handled in a culture favoring obedience (2) more complex environments should be better handled by a culture favoring autonomy and (3) a culture favoring autonomy should lead to higher robustness than one favoring higher obedience.

Formally, simulations are based on similar environmental setups than those of the baseline: performance in a simple yet intensive environment, performance when increasing complexity and finally, performance in bursting environment, with irregular income of tasks. New environments are variants of the ones from the baseline supporting longer tasks, but which can be split and performed concurrently by agents.

Results are presented in Figure 6.14, Figure 6.15 and Figure 6.16. They show similar properties with those explained in Section 6.4.1, with similar interpretations. The only difference concerns irregular environments. In these environments, culture promoting obedience tends to obtain higher performance when many tasks arise at once. Nevertheless, the gain in performance remains only marginal. In comparison, organizations in a culture promoting autonomy are able of handling more incoming tasks at once before dropping in performance, which is a clearer indicator of robustness.

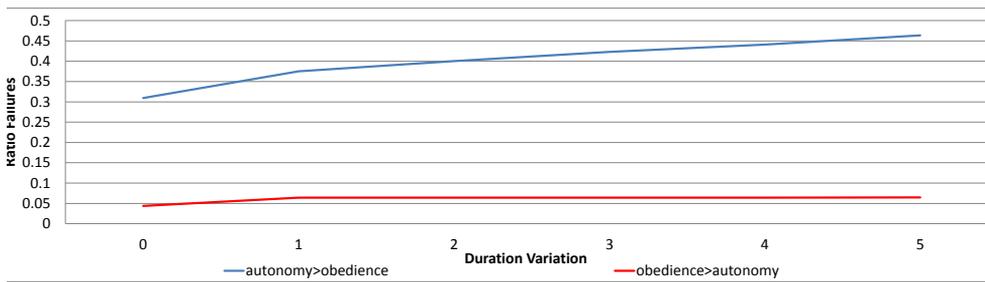


Figure 6.14 – Ratio of failed tasks (Y-axis) depending on the variability of the duration of incoming tasks (X-axis) in the intensive environment with the “cooperation” organization. Each curve represents a cultural preference towards autonomy vs. obedience. These curves indicate the organizational efficiency and flexibility in simple environments depending on culture.

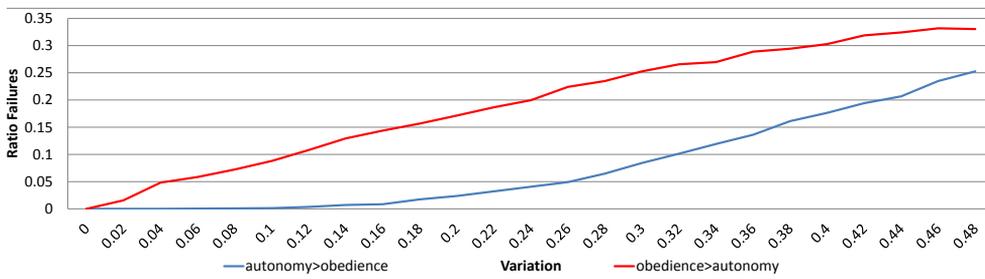


Figure 6.15 – Ratio of failed tasks (Y-axis) depending on the risk that incoming tasks are complex (X-axis) in the complex environment. Each curve represents a cultural preference towards autonomy vs. obedience. These curves indicate the organizational performance in complex environments depending on culture

6.8.3 Back to Theory

Recurrence

This simulation model illustrates that a culture supports recurrent interaction patterns. The interaction patterns observable in this simulation are the same as those observable in the baseline model.

More concretely, as for the baseline model, if obedience is culturally important then leaders manage the workload of subordinates and subordinates comply with orders of the leader. In this model, the decision power is centralized around the leader.

Conversely, if autonomy is culturally important, then subordinates manage their own workload. Leaders are simple communicators and have little information about subordinates. In this model, decomposing and exchanging tasks is decentralized to subordinates.

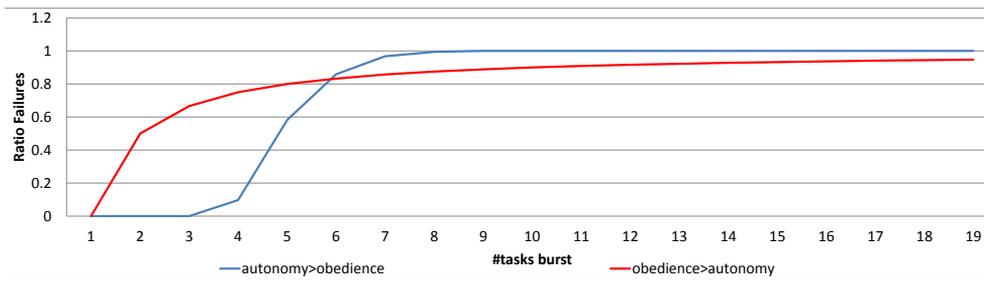


Figure 6.16 – Ratio of failed tasks (Y-axis) depending on the amount of tasks in each burst (X-axis) in the irregular environment. Each curve represents a cultural preference towards autonomy vs. obedience. These curves indicate the organizational robustness, depending on culture

Culture and Coherence

In this model shared cultures lead to similar coherent interactions as those from the baseline model. Likewise, similar incoherent interactions arise from cultural mismatches.

Culture and Collective Outcomes

This model illustrates the recurrent influence of culture on collective outcomes. As for the baseline model, culture supports similar abstract coordination patterns (centralization vs. decentralization). These similarities in terms of interactions lead to similar collective outcomes.

Coordination Mechanisms and Cultural Expressiveness

This model illustrates relationships between coordination mechanisms and cultural expressiveness by considering a situation that provides more choice to agents. Indeed, agents have to decide whether they split the task themselves or delegate it to another agent. Furthermore, subordinates can decide whether they delegate tasks to other agents. These additional choices are also motivated by culture. To that extent, this less strictly regulated coordination mechanism supports higher cultural expressiveness.

Conclusion

This section illustrates through a concrete example how culture leads to recurrent interaction patterns. In other words, culture supports the occurrence of similar abstract interaction patterns. These patterns occur relatively independently from a specific environment or coordination mechanism.

As a direct consequence, this model illustrates that interactions supported by culture can remain coherent in spite of a different context. Furthermore, this model shows that culture, as long as it can be expressed, has a relatively context-independent influence on collective outcomes. These two properties are particularly interesting when aiming at using culture for supporting coordination.

6.9 Further Evidence

At the start of this chapter, we introduced a general theory that explains how the influence of culture on individual decisions leads to the influence of culture on collective outcomes. Then, as for our decision architecture from Chapter 4, we need to get some confidence that this theory is (sufficiently) close to reality. This confidence is crucial for creating expectations that this theory makes sense and can thus be used in the following chapters and by other researchers.

As in Chapter 5, we need to consider first how much confidence we need. Indeed, as for other natural-science theories, our theory is based on (limited) observations. Thus, we cannot “prove” that our theory is correct. Nevertheless, through evidence, we can get sufficient confidence that this theory is accurate with regard to reality. More precisely, that this theory is accurate enough such that it can be relatively safely used in the following sections for achieving the goal of this thesis.

With regard to the goals of this thesis, we need to be relatively confident that our theory (1) matches the basic known properties of the influence of culture on coordination from human societies and (2) is applicable for multiple values. Here are five pieces of evidence that support the idea that our theory can be relatively confidently used.

First, our theory conceptually matches existing theories. Our theory assumes similar influences of culture (particularly, values) on individual decisions and collective outcomes as existing theories. Likewise, our theory matches available theories describing the interaction level (e.g. coherence by avoiding misunderstandings and mis-expectations).

Second, our theory matches culturally-specific evidence about the influence of culture on interactions. Some evidence is concretely displayed by the simulation presented in this chapter (e.g. centralization in high PDI-CD culture, decentralization in low PDI-CD culture, risk-taking in high MAS-CD culture, risk-avoidance in low MAS-CD culture). Additional evidence can be easily conceptually derived from descriptions of culture about individual and collective levels in considering interactions (e.g. uncertainty pushes to formalize interactions leading to higher efficiency but lower flexibility in high UAI-CD culture, adaptive interactions in low UAI-CD culture).

Third, our theory is generic. While being inspired by the link autonomy/obedience, our theory from Section 6.1 is conceptually independent of any specific set of values. The previous argument illustrates numerous examples of value-systems where our theory is useful for explaining the individual-to-collective link.

Fourth, the various concepts and link proposed by our theory are richly illustrated by the simulation proposed in this chapter. The simulation proposed in this chapter illustrates all the aspects proposed by our theory. In particular, this theory remained valid in numerous settings: different value-systems, matching and mismatching cultures, various organizations.

Fifth, our theory can be used for solving the seeming paradox we faced during the simulation of Chapter 5. Theories of culture relate culture to both context-dependent and context-independent influences on coordination. Our theory can be used for explaining both influences. We can explain the independence relationship with the recurrence of interaction patterns. For instance, in the simulation from Chapter 5, conservatism supports standardized resolutions. This form of resolution is more efficient when individuals have the adequate expertise (higher efficiency), but fails when complexity increases (lower

flexibility). With our theory we can also explain the dependence relationship between culture, coordination and performance in relying on cultural expressiveness. For example in our simulation, conservatism is expressed better than openness-to-change when the organization is a machine bureaucracies. In more details, both machine bureaucracies and conservatism prefer to rely on protocols, while this is not the case with openness-to-change. Similar explanation can be provided for understanding better why adhocracies match better openness-to-change than with conservatism.

Together, these pieces of evidence provide some good confidence that our theory makes sense with regard to reality. In particular, they suggest that, in spite of being originally inspired by the autonomy/obedience pairs of values, our theory is applicable for a wide range of values.

6.10 Chapter Summary

This section explains, through a theory and several illustrations based on simulations, how culture influences coordination, from its influence on individual decisions to its influence on collective outcomes. Overall, cultures can be seen as a drive pushing all agents in the same direction.

This drive pushes agents towards making decisions which, collectively, lead to the occurrence of a given abstract interaction pattern. The more a culture is shared (i.e. richer cultural models), the more this drive is precise and coherent, thus avoiding “cultural clashes” (e.g. mis-expectations, misinterpretations).

Abstract interaction patterns resulting from this drive are relatively independent with regard to a given decision situation. In various situations, similar interaction patterns are to be expected. Nevertheless, in spite of achieving similar interaction patterns, decision situations (e.g. coordination mechanism, environment) can promote or demote cultural expressiveness. This promotion or demotion depends on whether the situation supports rationally-acceptable alternatives that can be supported or not by culture (e.g. leadership supports the expressiveness of obedience). Thus, the decision situation, while not influencing the direction proposed by culture, influences how far culture influences a given decision.

What about the influence of culture on collective outcomes? Culture, by promoting context-independent interaction patterns, drives collective outcomes in a relatively predictable direction. This chapter investigates in particular the influence of mismatching culture on collective outcomes and leads us to unexpected results: cultural clashes can occasionally improve collective performance. While this observation is of great importance with regard to the next chapter of this thesis where we show the benefits of sharing a culture (why using culture for supporting coordination if mismatching culture can do as well), we also observe that better outcomes resulting from cultural clashes are relatively rare. In most of the cases, cultural clashes lead to performance loss, thus shared culture supports in general better outcomes.

Regarding the following chapter, this theory helps understanding how cultural values can support flexible coordination: by supporting coherent interactions. Through models of values, this theory can then be transposed within MAS and can be used for determining how cultures can be used as a tool for supporting for supporting flexible coordination

within MASSs.

Answering Research Question 1.b: How Does Culture (Indirectly) Influence Collective Outcomes Related to Coordination in Human Societies?

In this chapter, we explored further this research question. So far, existing theories implicitly suggest that sharing culture leads to better collective outcomes. The simulation provided in this section allows to further enrich this suggestion.

Sharing a culture generally but not necessarily supports coordination. Sharing the same culture generally supports better performance in situation of coordination, by avoiding misunderstandings and mis-expectations, more generally by creating *coherent interactions*. Nevertheless, societies sharing a culture do not necessarily strongly dominate (do better in any case) societies where no culture is shared. First, because a specific culture is inadequate for specific environments (e.g. obedience in too complex environments). Second, because sometimes (cultural) misunderstandings can actually lead to good performance “by accident”.

The other point we further explored in this chapter is the relative independence between culture and collective outcomes. Theories implicitly suggest that each culture supports a specific type of performance (e.g. value of obedience supports poor robustness, value of autonomy supports flexibility).

We acknowledge this independence relationship, but we moderate it. Culture appears to support a specific type of coordination performance because they support recurrent interaction patterns. For instance, in any environment, obedience supports centralization, which generally leads to poor robustness (bottlenecks, no redundancy).

Nevertheless, we moderate this “independence” by adding an aspect to take into consideration: cultural expressiveness. Depending on the situation, culture may be more or less easy to express (e.g. the influence of obedience is limited in flat organizations). To that extent, while performance is driven by culture in a direction that is problem-independent, the “strength” of this drive is problem-dependent.

Answering Research Question 1.c: How Does the Influence of Culture on Individual Decisions Lead to the Influence of Culture on Collective Outcomes in Human Societies?

This question is the core target of this chapter and crucial for answering our end-goal research question.

In order to tackle this question, we introduce an intermediate layer for considering this problem: interactions. Interactions are close enough to the perspective of individuals for being easily explained in terms of individual decisions and close enough to collective decisions for explaining the latter from interactions.

From individual decisions to interactions. Culture influences individual decisions, amongst others, through values. Values indicate which alternatives can be expected to be “good” when making decisions. When individuals share the same culture, then culture drives specific types of interaction patterns. In general, values supported by culture are

abstract, leading to abstract individual drives and abstract interaction patterns. These individual drives can be implemented in many situations, leading to recurrent interaction patterns. Finally, values supported by culture tend to drive towards decisions that can be considered as desirable for individuals (e.g. stimulation, self-discipline) or for the collectivity (e.g. self-sacrifice, generosity). Furthermore, values supported by culture tend to support coherent interactions. They allow creating expectations about decisions of other individuals and decisions supported by cultural values tend to collectively “match” each other.

From interactions to collective outcomes. These abstract, coherent and recurrent interaction patterns can be easily related to what is known about collective outcomes. In particular, the recurrence of interactions explains why a specific culture supports relatively context-independent collective outcomes. The coherence within interactions supported by values explains why sharing the same culture tends to support collective outcomes.

More generally, explaining collective cultural phenomena based on the cultural influence on individual decisions is greatly simplified in considering the interaction level: first in connecting individual decisions to interactions and interactions to collective outcomes. This perspective of understanding cultural phenomena was hinted to us by the MAS methodology. This perspective offers a relatively new way for understanding how culture influences societies which appears to be particularly promising for using it in order to coordinate MASs.

Answering Research Question 2: *How to Model the Influence of Culture on Coordination?*

This chapter provides another example for modeling the influence of culture in a situation of coordination. In this chapter, we modeled values in designing value-compliant decision rules.

Values can be used implicitly, as a guideline for designing more concrete decision-rules. In this case, values can be used for justifying why using a decision rule instead of another (e.g. subordinates blindly accept orders because they are obedient. No need for deeper reasoning).

This model is particularly adequate for designing social simulations. One of the key principles for building social simulations that are explainable, controllable and computationally manageable consists in keeping agent decisions as simple as possible. The need for accuracy at the individual level is less important since the focus is on collective phenomena. In this context, simple rules are particularly adequate for modeling decisions. These rules can then be externally justified with regard to underlying more complex decision processes.

In conclusion, values can be *implicitly* used for justifying the implementation of value-compliant simple decision rules.

7

Using Value-Systems for Supporting Coordination in Artificial Societies in Practical System Design

Chapter Summary

This chapter investigates how system designers can use human-like culture, more specifically value systems, as a tool for supporting coordination in artificial societies. First, we investigate the type of coordination-problems for which value-systems can be used for supporting coordination with regard to other available techniques. Second, we investigate practical considerations to be considered when implementing this solution.

Based on the answers provided by previous chapters, this final chapter provides an answer for Research Question 3 raised at the start of this thesis: *How to use models of human culture as a practical tool for supporting coordination in artificial societies?* Previous sections, by highlighting the relative appropriateness of value-systems, let us refine this question into how *value systems* can be used as a practical tool for coordinating artificial agents. We acknowledge that others aspects of culture can be useful for supporting coordination (e.g. rituals), but these aspects are less appropriate with regard to our goal, as justified in Section 4.1.1.

The word *practical* used in this question conveys a crucial meaning. As further depicted in Section 2.1.3, we consider the concrete perspective of system designers and their inherent challenges. These challenges are numerous, such as bounded means for coordinating agents, possible imperfect information about the problem and the environment, need to achieve a certain level of performance. Even deciding how to coordinate agents is costly.

In order to decide efficiently about how to coordinate, system designers generally make decisions on two levels. First, system designers make global or abstract decisions about how to coordinate, such as selecting the abstract coordination techniques they want to use. Then, system designers make more specific decisions for implementing their abstract coordination decisions (e.g. implementing the selected coordination technique).

In this chapter, we investigate solutions for supporting decisions about these two levels when using value-systems, based on the work of Vanhée et al. (2014b). In order to support the decision about the first level, we investigate in which cases value-systems are appropriate for supporting coordination. This investigation is performed through Research Question 3.a: *When are value systems appropriate for supporting coordination?* In order to support decisions about the more detailed level, we investigate Research Question 3.b.: *What are the practical considerations for designing value systems to support flexible coordination?* Throughout this investigation, we also highlight design challenges that can be raised by practical considerations and introduce several techniques for handling these challenges.

Regarding other chapters, the current final chapter relies on the work developed by all previous chapters of this thesis. Chapter 4 and Chapter 5 show that the influence of culture on coordination can be operationally replicated within MASs. Chapter 2, Chapter 3 and Chapter 6 introduce a theoretical and conceptual background, for understanding how values influence coordination in human societies. This background offers strong guidelines for determining the consequences to be expected by the integration of cultural values on coordination within MASs. Based on these previous chapters, the current chapter can further elaborate about how the influence of culture on human-like values can be used as a tool for supporting coordination within MASs.

The structure of this chapter is as follows. Section 7.1 investigates the type of coordination problems that are well handled by value systems. Then, Section 7.2 further investigates practical considerations for designing value systems for supporting coordination and their integration within agent decision processes. Finally, a conclusive section gathers the answer to both questions for answering our end-goal research question.

7.1 Answering Research Question 3.a: *When are Value-Systems Appropriate for Supporting Coordination?*

This section investigates the range of coordination problems for which value-systems are adequate for supporting coordination. In particular, we consider a practical setting, where coordinating and making decisions about coordinating is costly. Thus, in order to be used for a given coordination problem, value-systems should not only be applicable, but should also add to other coordination techniques.

Given the system-designer-oriented perspective taken by this thesis, we focus on providing *efficient* rules (time efficient and discriminative) for deciding whether to use value-systems or not. This need for efficiency prevents completeness, because completeness implies checks that are too costly (e.g. need to consider whether technical details introduced in Section 7.2 can be performed). Instead, we aim at providing basic general rules for easily deciding whether value-systems can be appropriate or not.

Two main general conditions are to be met for appropriately using value-systems. The first condition is *theoretical requirements*. These requirements are basic preconditions drawn from theory that make that, if they are not met, value-systems can hardly be used for supporting coordination. The second condition is the *relative benefits* of using value-systems. These benefits indicate the range of coordination problems where value-systems can *add* to the support provided by other coordination mechanisms.

Theoretical requirements for using value systems are introduced in Section 7.1.1. Then, relative benefits from value systems for supporting coordination are introduced in Section 7.1.2.

7.1.1 Theoretical Requirements for Applying Value-Systems

This section introduces the core conditions for determining whether, from a theoretical perspective, value-systems can be used for supporting coordination. These conditions are derived based on the theories depicting the influence of values, introduced in Section 4.1.2 and our explorations in Chapter 6.

In order to use value-systems for supporting coordination, agents need to face (at least some) decision situations that respect all of those conditions. These conditions aim at raising decisions (time efficient and discriminative) can both influence coordination and be influenced by value-systems. First, the decision of the agent should influence (directly or not, knowingly or not) an interaction. Second, agents should have enough “freedom” when making this decision. Third, value-systems should support alternatives of this decision. Fourth, this decision should be difficult to handle with another decision mechanism.

Situation of Interaction

In order to influence coordination, value-systems should influence decisions that influence interactions. Indeed, as introduced in Section 2.1.1, coordination consists in “harmoniously” solving situations of interactions. Thus, in order to influence coordination, value systems should influence decisions in situations of interactions.

Note that agents are not required to be aware of being in a situation of interaction for being coordinated by value-systems. Likewise, value-systems may influence decisions without interaction. For instance, “timeliness” can support coordination (e.g. agents deliver tasks on time, making that schedules are reliable). Nevertheless, agents influenced by this value do not need to explicitly consider that they are in interaction for appropriately coordinating (e.g. expect that the tool for operating will be available without delays).

Freedom for Making Decisions

In order to be influential, value-systems should influence decisions with multiple alternatives. In some sense, agents need to have some “freedom” about which outcome to select (no obvious outcome, no overwhelming coordination constraints). In practice, agent’s freedom can be easily restricted by environments or other coordination techniques (e.g. regimented norms), as illustrated in Section 6.7. Restrictions on possible agent’s freedom restricts in turn the potential support which can be provided by value-systems.

As a general rule, the more agents are given freedom, the more values can influence decisions and thus coordination.

Values Can Discriminate Alternatives of These Decisions

In order to influence decisions, values should discriminate the various alternatives of these decisions. Otherwise, all outcomes are “equally good” according to values, nullifying the influence of values towards one outcome or another.

In general, this problem is solved when decisions supported by values are connected to a rich conceptual background. Indeed, this background is crucial for providing the necessary ground for performing (influential) evaluations (e.g. “good” or “bad” instead of “irrelevant”). These concepts can be provided both by the environment (e.g. challenging tasks, uncertainty) and by other coordination mechanisms (e.g. leader and subordinates).

As an illustration inspired by van der Weide (2011), “select house to rent” can be related to numerous concepts (e.g. space, distance to work, distance to city center, cost of the rent). These concepts can, in turn, be related to values (e.g. comfort, autonomy, security). A decision such as “select which foot to start walking with” provides much less background for connecting values (e.g. in armies, this decision is better supported by rules).

Value-Systems Should Not Be Taken Over by Other Decision Mechanisms

Value-systems can be taken over by other decision processes. In particular, rational decision processes can and should take over (irrational) value-systems when possible. In order to be useful, value-systems should be used for decision situations that cannot (timely) be handled by rational processes (e.g. by providing heuristic rational evaluations, arguments for selecting rational-enough outcomes).

The range of decision situations that are hard to handle by rational decision processes can be characterized by several aspects such as conceptual complexity, computational complexity, uncertainty, interactions with others, too many possible alternatives or abstraction.

Conclusion

This section introduces a set of conditions for using values. In conclusion, value-systems can be used for supporting coordination when agents face conceptually-rich decisions in situations of interactions that are hard to solve rationally. Ideally, these decisions raise “rational social dilemmas”, which can be connected to values (e.g. an opportunity for saving time arises, but may impact on my colleagues. Shall I take the initiative by myself? Shall I ask others? Shall I ask my boss?).

As a general rule, these conditions are met in environments raising complex decisions or a wide range of decisions, where value-systems can provide heuristic guidelines can add to rational choice. These conditions can also be met in simpler environments if abstract longer-term decisions are to be made (e.g. the agent has multiple approaches for tackling a problem).

7.1.2 Relative Benefits: Specificities of the Support Provided by Value-Systems

This section introduces the general conditions for which value-systems add to the support provided by other coordination techniques. Value-systems are particularly adequate for coordination problems meeting at least one of those conditions.

The Problem Involves Handling a Wide Range of Decisions

Value systems are particularly appropriate for supporting coordination for a wide range of coordination situations (e.g. many types of interactions, evolving environments). Indeed, as highlighted in Section 4.1.2, values are generally relatively abstract and general. Thus, they can be applied for providing a principled support for a broad range of decisions. They are even relatively appropriate for handling unexpected decision because their support is relatively abstract, relying on relatively limited *a priori* expectations.

For instance, the value of “obedience” can be abstractly used for supporting a broad range of decisions: does this decision go against orders? Is the leader aware of this? Shall I decide this by myself or report? This abstract guideline can then be implemented for concrete decisions (e.g. is this action going against my orders?). If an unexpected situation is encountered, value-systems provide a range of basic and generic aspects to take into consideration that are crucial for maintaining coordination.

In comparison, more traditional coordination mechanisms (e.g. protocols, norms, organizations, stygmergy, POMDP-like) are not appropriate for handling a wide range of situations. Indeed, the complexity of designing these mechanisms increases very fast with the range of situations to handle. Furthermore, these coordination mechanism generally require strong *a priori* expectations. Thus, such coordination mechanism can easily fail if the range of situations has not been handled properly by system designers (e.g. unexpected situation, failure of an action, new goal). On the contrary, value-systems offer a general “backup” for supporting decision and coordination when other coordination techniques failed.

There are more flexible coordination mechanisms (e.g. dynamic norm setting, abstract organizations, cooperation) that are more appropriate for handling a wide range of decisions. Value-systems add to these techniques by providing another perspective for tackling underlying design complexity. Indeed, value-systems provide generic coordination-supporting principles for driving individual decisions. This adds to fixed interaction patterns (cooperation), interaction structures (abstract organizations) or restrictions (abstract norms).

The Problem Involves Complex Decisions

Some problems raise decisions situations that can be inherently complex. This complexity may be the conceptual complexity raised by the domain (e.g. many variables to take into consideration, hard to determine all possible influences) or computational (e.g. need to look ahead in the future, consider many possible plans).

Coordinating agents acting in such complex environments is relatively difficult to handle with traditional coordination techniques. These techniques consists in letting the system designer deciding *a priori* which alternative shall be picked if a decision can influence coordination. By having to be designed *a priori*, these techniques require from system designers to handle all possible situations beforehand. Thus, when environmental complexity increases these techniques quickly lead to high expenses for handling all possible cases.

More flexible coordination techniques, allow partly overcoming this issue by letting agents coordinate on the fly. Nevertheless, these techniques generally rely on rational decision processes for making decisions. However, these processes are quickly limited due to

inherent decision complexity. Heuristics can be used to this end, but unless being similar to values they may not be appropriate for supporting coordination. In addition, complexity limits possible communication about coordination (e.g. too many things to discuss: information about the environment, plans, problems, goals, concepts and so on).

Value-systems add to this by providing a rich framework for making decisions and communicating about these decisions, as shown by van der Weide (2011). Indeed, values provide general guidelines and expectations for making coordination-supporting decisions that are relatively less sensitive to high complexity. For instance, if timeliness is important, agents can easily use these values for preferring timely plans, which are “good”. This value can be used for creating expectations that tasks of others will be done on time. Likewise, this value can be used for indicating that a plan is “good” because it allows to be on time. All of these supports for making decisions are applicable relatively independently from underlying complexity.

System Designers Need Support for Achieving a Precise Type of Performance

Value-systems are relevant for driving coordination towards a given performance profile. First, because various value-systems support a range variety of performance profiles (e.g. obedience with efficiency for simple problems, autonomy for robustness). This property is illustrated in Chapter 3 that relates the various cultural profiles with various influences on performance. Second, because the influence of value-systems on performance is relatively independent on a given problem, as explored in Chapter 6. Thus, value-systems provide strong *a priori* indications about the type of performance they support. Together, these properties highlight that value-systems can be used for supporting a wide range of coordination performance for a large range of environments.

This property is relatively unique in comparison with other techniques. As highlighted in Section 2.1.3, other techniques tend to support a specific type of performance profile (e.g. rules tend to be more appropriate for efficiency, cooperation for flexibility). When putting in practice these coordination techniques, system designers can influence the direction given to the performance profile (e.g. rules of redundancy for increasing robustness). Nevertheless, other techniques tend to be more environment-dependent and thus, the actual direction of a technique can be difficult to determine before considering implementation details. Value-systems add to these techniques by offering a wider range of potential performance profiles and relatively clear indication of the support before implementation.

The System Designer Wants a General Design Guideline

Value-systems add to other coordination techniques by concisely and meaningfully introducing general design guidelines for the whole system. In some sense, value-systems indicate the general *social paradigm* (or mood, or mindset, or atmosphere) to be expected in the system. Shall it be “all for the hierarchy” or “free for all”? Shall we make sure that things go fast through tight rules or be more relaxed and adaptive?

These guidelines are particularly crucial both for system and agent designers. These abstract guidelines indicate what is important and to focus on when coordinating agents and when being coordinated (e.g. rules, objectives, social structure). In some sense,

other coordination approaches can be seen as “practices” (in the sense presented in Section 2.2.3), which implement this abstract direction. Thus, value-systems provide a general direction for coordinating agents, precisely like values in human societies.

For agent designers, this guideline is useful for briefly introducing the social paradigm on which agents are based. This social paradigm indicate to agents what is important and “good” when acting and interacting. In addition, this guideline is crucial for compactly providing general expectations about society and the environment. For instance, “obedience” is important, leader agents can strongly expect sharp compliance from subordinates. This aspect is particularly relevant when multiple designers are in charge of designing agents. Indeed, value-systems help designers to avoid personal “social paradigms” to influence their agents.

In comparison, other coordination techniques do not introduce this abstract and concise social paradigm. In most of the cases, coordination techniques introduce behavior-oriented or outcome-oriented directions. Nevertheless, if the “spirit of the law” can sometimes be difficult to derive. This aspect is crucial for coordinating agents: relying on different social paradigms about interactions can easily lead to disappointments in spite of relying on well-working coordination mechanisms. Hofstede et al. (2010b) illustrate this in showing problems that arise when putting in interaction (using a negotiation protocol) agents with different (culturally-driven) mindsets.

Conclusion

This section highlights how value-systems can add to other coordination techniques. Value-systems are particularly appropriate for supporting coordination when agents face a wide range of complex decisions. Thus, value-systems offer a particularly useful solution for achieving flexibility and robustness which is one of the core pitfalls of available solutions.

Furthermore, value-systems can add to other coordination techniques by providing an abstract guidance, a “mind-set”, both for designing and coordinating agents. This guidance is particularly interesting because value-systems can relatively generically support the occurrence of a wide range of performance profiles.

The properties for considering coordination, introduced in Section 2.1.3, can be considered with value-systems. Value-systems are an explicit form of coordination: value-systems can be given as an external coordination constraint to agent designers. Value-systems are (particularly) adequate for coordinating in open systems. Value-systems do not impose constraints on goals set by agents, even if they influence mechanisms of goal adoption. They partly rely on the social configuration but do not require any social structure. They influence the type of social and environmental structures that can emerge (e.g. hierarchies if obedience is important). The locus of control is local, social relationships are known locally. No central entity is required for managing interactions. This form of coordination is also scalable (even though, some theoretical research, such as Beck and Cowan (1996) indicate that various cultures may better match various social sizes).

7.1.3 Conclusion: Pragmatically Deciding When to Use Value-Systems for Supporting Coordination

Value-systems can be considered for supporting coordination in a wide range of situations.

The most straightforward use of value-systems consists in tackling complex, dynamic and uncertain environments and (even in simpler environments) for making abstract decisions in situations of interaction. These environments match the requirement for using value-systems. They offer conceptually-rich decisions, when relying on rationality is expensive in terms of computational or design costs. They can involve situations of interactions that are difficult to handle with more traditional coordination techniques. Thus, these situations require that interaction-related decisions are made by agents. In addition, these environments and decisions correspond to those where value-systems add the most to other coordination techniques. These environments require making a wide range of decisions and can raise unexpected failures. These problems are relatively well handled by value-systems that provide a general abstract “backup”. Furthermore, value-systems provide abstract coordination-supporting decision heuristics that remain applicable in spite of decision complexity.

For similar reasons, value-systems can be used for supporting system flexibility and robustness. Value-systems provide a generic support for coordinating agents. This generic support can then be implemented on the fly for achieving a wide range of goals (flexibility) or as a backup, for supporting coordination when other coordination techniques failed. The support for flexibility and robustness adds particularly to existing coordination techniques (e.g. norms, protocols, organizations) that are either not ideal for handling these aspects or by offering a cross-perspective.

Last but not least, value-systems are adequate for providing general design guidelines. They indicate the “social atmosphere” to be found within the system. This atmosphere can be used for streamlining the design of agents, but also more specific coordination mechanisms. In addition, these various “social atmospheres” can be related to generic performance profiles. Thus, by picking a value-system, good expectations can be made about the overall performance of the system.

7.2 Answering Research Question 3.b: *What Are the Practical Considerations for Designing a Culture to Support Flexible Coordination?*

This section investigates practical aspects to be considered and solutions for designing value systems in order to support coordination.

In order to use value systems for supporting coordination, system designers have to determine several elements. System designers have to define the set of values to be used, combine them in order to determine the value system and finally provide them to agent designers that have to integrate them within agent decision processes.

Each following subsection introduces one of these elements. In Section 7.2.1 we investigate what are the desirable aspects for modeling coordination-supporting values. In Section 7.2.2 we investigate how to select the values to integrate within a value-system.

Then, in Section 7.2.3 we further describe how to combine values together in order to create value-systems. Finally, in Section 7.2.4 we show how to integrate value-systems within agent decision processes.

7.2.1 Practical Aspects to Consider When Designing Values

Values constitute the basic independent components of value-systems. Thus, as a first step for designing value-systems, we need to consider which values that are to be picked. Several values have been identified in the literature (e.g. Schwartz values Schwartz (1992)), but the list is not guaranteed to be exhaustive.

In this section, we introduce core properties for modeling human-inspired coordination-supporting values. Since basic properties for modeling human-like values were introduced in Chapter 4, we assume that these properties are met and we focus on properties for obtaining coordination-supporting values.

We identify two core properties. First, values should be applicable for supporting decisions. Second, values should be oriented towards coordination.

In the following sections, we introduce each of these properties. Finally, we introduce examples of coordination-supporting values.

Values Should Be Applicable for Supporting Coordination-Influencing Decisions

As highlighted in the requirements, described in Section 7.1.1, using a value for supporting coordination makes sense under certain conditions: agents should need an irrational (efficient) support for making coordination-influencing decisions.

As a general solution, for handling this practical issue, we propose to rely on *abstract* values. While there is no theoretical need for abstract values, as shown by Miceli and Castelfranchi (1989) (e.g. brushing teeth can be a value), values can (and tend to) be abstract in human societies, as shown by Schwartz (1992). From an implementation perspective, the more a value is abstract, the more this value can be used for supporting coordination. As a downside, too abstract values may be hard to connect to very specific expectations (e.g. obedience does not tell exactly how a business is run). Nevertheless, the aim of using value-systems is not to provide such specific expectations, but a general direction.

Techniques for tackling this drawback are proposed in the following section. Section 7.2.4 introduces a technique for making abstract values more concrete (e.g. implementing “mastery” through more concrete values of “mastering speed” and “mastering direction”). In Section 8.3.3, we show as future research the possibility to derive specific practices from values (e.g. obedience supports frequent reports).

Values Should Be Oriented Towards Supporting Coordination

In addition to being expressible for supporting coordination, a value still needs to promote coordination-supporting decisions in order to be useful in our case.

A first practical concern to consider is that an appropriate value should help discriminating the various outcomes of coordination-influencing decisions. Otherwise, the addition of this value for supporting coordination is null. For instance, “honesty” is pointless

if no lies can be made. This issue can be handled by considering the type of concepts that coordination-influencing decisions should raise.

A second practical concern to consider is the appropriateness of these values for supporting coordination. Some values can hardly be useful for supporting coordination (e.g. "hedonism"). In the given state of our knowledge, the best advice for handling this issue we can propose consists in relying on human-inspired values. In particular, values promoted by culture are generally appropriate for supporting coordination. As indicated in Section 4.1.2, values supported by culture do necessarily support coordination (e.g. self-punishment, proselytism). However, they generally drive towards coordination-supporting decisions, particularly when they are shared (e.g. obedience supports coordination if both leaders and subordinates share this value). To that extent, cultural values tend to particularly influence and support coherent coordination patterns, as theorized and illustrated in Chapter 6.

Examples of Values

This section introduces examples of values that support coordination in human societies. These values were used throughout the research-work related to this thesis for building models and theories. Most of these values can be related to HCDs, which were a prime source of inspiration for influential values. Nevertheless, more values that are adequate for supporting coordination can be found in the literature.

PDI-CD-Related Values (e.g.: social order, obedience and autonomy): These values influence the repartition of informal decision power. These values are applicable as long as agents are tied with power relationships and when power is not completely formally defined. Their influence is relatively abstract and can thus be concretely implemented in numerous situations. They discriminate alternatives of individual decisions and clearly support coordination (implicit collective drive towards (de)centralization). More information about the influence of these values can be found in Chapter 6.

MAS-CD-Related Values (e.g. assertiveness, mastery, harmony and modesty): These values influence how agents manage risk-taking decisions and failures as depicted in Section 3.1.3. Their influence is abstract and can be made concrete for a wide range of situations. They provide appropriate decision heuristics when they are shared (e.g. provide accurate information about how to interpret proposals of others). They clearly support coordination (e.g. balancing workload, creating expectations about risks of failures from other agents). They discriminate alternatives of individual decisions. More information about the influence of these values can be found in Chapter 6.

UAI-CD-Related Values (e.g. security, stability, creativity, adaptiveness, self-direction): These values influence how agents react to uncertainty and norms as depicted in Section 3.1.4. These decisions are abstract and applicable in a wide range of situations. These values support coordination when they are shared (e.g. maintaining low uncertainty through complying with norms vs. individually managing uncertainty and make adaptive decisions). They discriminate alternatives of individual decisions (e.g. following norms). They support two types of coordination: either maintaining low uncertainty through standardized behaviors (other can thus rely on strong expectations) or assuming uncertainty and handling each task adaptively (other can thus adequately expect uncertain outcomes).

Sequential versus Synchronous Perception of Time: The influence of these values is inspired by Hampden-Turner and Trompenaars (1993) who describe two paradigms when considering time management: sequential and synchronous. In considering time as sequential, individuals think about time as a strict sequence of events to be respected. Conversely, in considering time as synchronous, time is a resource to optimize (e.g. maximizing activity). Both values are abstract and can be applied in a wide range of situations. Both support coordination (tight schedule vs. opportunity-taking and delays). They indicate what is “inherently good” when shared (e.g. being in time for a meeting is good when others are also in time; optimizing time is good when delays are frequent). They discriminate alternatives of individual decisions.

Conclusion

This section highlights several practical aspects for considering what to use as values and guidelines for handling these practical aspects. Criteria should both respect core properties of values (heuristic for making irrational decisions) and be adequate for supporting coordination (support actual decisions to be made, lead to better interaction patterns).

7.2.2 Practical Aspects for Choosing the Values to Be Integrated

The second step for building value-systems consists in determining the *set of values* that should be integrated within the system. This set of values is then to be put in relation and implemented within agent decision process.

This step raises two core practical concerns. First, values should together be adequate for supporting coordination. Second, integrating values strongly influence design costs for the system, raising the question of limiting these costs.

Selecting Coordination-Supporting Sets of Values

Designing adequate values as presented in Section 7.2.1, is necessary, but not sufficient for designing coordination-supporting sets of values. Indeed, the influence of multiple values can enter in interaction and lead to unexpected results. To that extent, multiple practical aspects have to be considered when selecting the set of values to integrate within the system.

The Set of Values Should Cover Coordination-Influencing Decisions: Unlike individual values, which can cover only a part of coordination-influencing decisions, these decisions should be sufficiently covered by the set of values. The union of the coordination-influencing decision for each value from the set should cover enough coordination-influencing decisions in general.

Given a set of values, system designers should consider the set of (un)covered coordination-influencing decisions. If this set of covered decisions is insufficient, system designers should consider cover additional coordination-supporting decisions, either through adding additional values or another coordination mechanism.

Coordination-Influencing Decisions Supported by Values from the Set of Values Should Not Overlap (Too Much): Each value supports a range of decisions and a decision can be supported by multiple values. Thus, multiple values can interact at the level of a decision. Can this interaction be beneficial for supporting coordination?

First, multiple values may contradict each other. For instance, different alternatives for the decision of “lying for covering up one’s leader” are supported by different values. “Yes” is good from the perspective of obedience and “No” is good from the perspective of honesty. This case is not ideal for supporting coordination: these conflicts weaken the support provided by values.

Second, multiple values may agree with each other. For instance, the answer “yes” to the decision “telling the truth for covering up one’s leader” is supported both by honesty and obedience. In this case, both values are not useful for supporting coordination either: one was sufficient. Nevertheless, more values can sometimes be useful for sharpening the support provided by values and thus making quicker decisions.

In conclusion, overlapping values are not useful for supporting coordination *per se*. Nevertheless, this overlap is (almost) inevitable due to the abstract and generic support provided by values. This overlap is still useful for efficiently tackling a wide range of situations. For instance, obedience and avoiding danger can enter in conflict. They induce a gray area for situations where “danger equals obedience”, but they provide sharp support for the many cases where one aspect takes over the other.

As a general guideline for tackling this issue, we recommend selecting values that can be conceptually complementary. The redundancy between the drive of all these values on decisions should be limited. For instance, maybe only two of the three values of “social order”, “discipline” and “rule-compliance” are sufficient for supporting decisions that are to be made. The last value would overlap with others without providing further support. Second, by avoiding introducing values that enter too much in contradiction with other values. For instance, a system highlighting the cultural importance of both autonomy and obedience is likely to raise moral dilemmas that will not help agents to make decisions or create expectations about others and thus fail to support coordination.

Values Should Support Desirable Interaction Patterns: In Chapter 6, we explained that cultural values drive the occurrence of abstract coherent interaction patterns that are useful for supporting coordination in general. When considering multiple values, we can expect the merger of the multiple abstract interaction patterns supported by each value. This merger can be more or less adequate for supporting coordination. For instance, security and assertiveness may limit the expression of each other’s interaction patterns: formalized interactions supported by security are likely to conflict with risk-taking decisions supported by assertiveness.

In addition, the interaction pattern supported by a set of values has to be tuned for the coordination problem. For instance, centralized interactions supported by obedience are appropriate for supporting coordination for relatively simple problems, but are less appropriate for more complex problems or when seeking robustness.

Managing Design Costs

The other core aspect to consider when picking values are design costs. More values increase design costs, because each value has to be integrated within agent decision processes. Furthermore, more values imply more possible conflicting overlap between values that has to be handled (e.g. indicating the relative importance between values, additional reasoning mechanisms for agents).

These costs lead to the typical trade-off between coordination support and design

costs. As a general principle, we can suggest to limit the number of values to pick, select values that are easy to evaluate and avoid conflicting values.

Conclusion

This section introduces core practical aspects for determining the set of values to integrate within a value-system. Values should drive agents towards making decisions leading to interaction patterns that support coordination. Values should match out overlapping too much and support coordination-influencing decisions to be made. Furthermore, the set of values to integrate within the system strongly influences the overall cost for implementing the value-system.

7.2.3 Practical Considerations for Combining Values

As presented in Section 4.2.1, values can easily conflict for some decisions. Several solutions are available for handling these conflicts. For instance, if obedience and timeliness are chosen, in which situations subordinates may disobey their leaders in order to be on time? How much delay can be allowed? For which orders? A key mechanism consists in (relatively) ordering the importance of various values (e.g. obedience is more important than timeliness). This mechanism supports *consistent* behaviors, which is crucial for creating strong expectations about the behavior of an agent. Culture plays an important role when fine-tuning this mechanism by promoting a common value-order.

In our setting, sharing a common value-system is crucial for creating expectations, both for other agents and system designers, and thus for supporting coordination. Nevertheless, determining the relative order between values can be difficult to decide and to implement. This section introduces a set of practical considerations when combining values.

Internal Coherence

As introduced in Chapter 4, agents should try to maximize the satisfaction of all their values. They should only select “pareto-optimal” outcomes according to values (not selecting an alternative if another alternative is equivalent or better in all points). In our model of values, we expressed this property through the “*ceteris paribus*” (Brafman and Domshlak (2013)) relationship.

Avoiding Absolutism

A simple solution for combining value-systems consists in strictly ordering values using a lexicographical order. For instance, maximize obedience. If two alternatives have the same level of obedience, then maximize timeliness.

In our experience, these value-systems lead to extreme behaviors that are not appropriate for designing flexible or robust systems. These value-systems easily lead to extreme behaviors (e.g. obeying orders even if it incurs long delays for others) and extreme variations in behaviors (e.g. as soon as the order cannot be obeyed, then be in time for any interactions with other).

Since flexibility and robustness are the core benefits of using value-systems, such absolute value-systems are to be avoided. They are more easily represented with rules for instance.

Design Solution: Relative Importance

As a general solution that we used in Chapter 4 and in Chapter 6, we propose to relate values with a relative importance. For instance, “obedience is slightly more important than timeliness” or “obedience is dramatically more important than timeliness”. This solution is adequate for avoiding too extreme behaviors while maintaining internal consistency. Furthermore, this solution maintains a form of “global consistency”: achieving any (positive) degree of satisfaction for a value will be preferred against achieving the same degree of satisfaction of less important values. Finally, this solution is relatively inexpensive in terms of design costs while still being adequate of supporting a wide range of resolution of conflicts between values.

Conclusion

In this section, we investigate basic practical concerns for designing value systems in order to better handle situations where values conflict with each other. In addition, we propose a simple and cheap solution for ordering values.

This value-system can be seen as a specification that system-designers provide to agent-designers. This value-system raises core values that agents should consider as important and behave in accordance with. The next step consists in integrating these values within the agent decision process.

7.2.4 Practical Concerns for Integrating Value-Systems within Agent Decision Process

Once the value system is designed, the last step consists in integrating this value system within the agent decision process. This step is crucial since it determines how values will concretely impact agents.

Concretely and operationally conceptualizing the influence of values on decisions is a challenging task. We proposed such a value-sensitive model in Chapter 4. This model is sufficient for replicating the core properties we needed in simulations: obtaining decision outcomes that are coherent with theories of values (agents try to maximize the satisfaction of their values). Nevertheless, many more cognitive aspects of values can be added when implementing our general decision architecture, as indicated by Miceli and Castelfranchi (1989). In particular, values can be used for “shortcutting” exploration rational decision processes, by telling that a branch is “good” or “bad”. Such cognitive additions can be particularly relevant in practical decision problems.

From a modeling perspective, the range of value-sensitive agent decision models remains limited at the moment. So far, the best model available is provided by van der Weide (2011). While this framework provides some basic insight about how to obtain value-coherent decisions. Nevertheless, a lot of work remains to be done in this direction for integrating rich cognitive aspects as depicted by Miceli and Castelfranchi (1989).

In this section, we introduce several core practical aspects to consider when concretely modeling values in order to support solutions and provide insight about solutions for handling them. Most of this insight is based on the research of Miceli and Castelfranchi (1989); van der Weide (2011).

Representing and Value Systems

A first aspect to take into consideration is the way agents internally represent value systems. This representation impact how values influence decisions and thus when values can be used. We can consider three main solutions for representing and using value-systems for making decisions.

The most direct solution consists in *implicitly* integrating value-systems within the agent decision architecture. Value-systems are used as a guidance or a justification when designing agent decision processes. Values are not explicitly introduced within agents for making decisions. To make a parallel with norms, agent designers would try to have agents behaving according to norms without having agents directly manipulating norm-related concepts (e.g. making the car drive on the right side of the road without integrating the road code within the agent decision process). This solution is particularly appropriate for designing social simulation. We used this representation in Chapter 6.

The second solution is based on frameworks *à la* van der Weide (2011). This solution consists in *explicitly* integrating values within agent decision architecture. This solution is the one we implemented in our simulations. In this case, agent decision processes are composed of two parts. First, agents generate a set of rationally-similar alternatives. Then, the various alternatives are evaluated by each value, generating tuples of evaluation. These tuples are compared with each other with the value system. The more satisfactory tuple indicates the decision to be selected.

A third solution, based on the work of Miceli and Castelfranchi (1989) consists in integrating values even deeper within the decision architecture. Values are entangled with objective rational evaluations. In this setting, values are used as heuristics for stopping on or pointing “good” solutions to rational reasoning.

Selecting the Right Decisions to Support with Values

A second aspect to consider is which decisions can be supported by values. As highlighted in Section 7.1.1, not all decisions are appropriate for supporting coordination.

From the agent perspective, values are particularly appropriate for supporting decisions that are recognized to be *impactful*. While no extensive definition of impactful decisions are available, these decisions can be characterized by the range of decisions that imply a form of commitment. This commitment can be characterized by expected difficulties for revising the consequences of this decision once made (e.g. costs, impossibility) or decisions with relatively expected important influences on individual welfare or on others. For instance, the decision “changing job” is generally more impactful than “drinking a tea or a coffee” in a normal setting. Norm-related decisions (e.g. adopting or violating a norm) tend to be impactful in general. In this track, *abstract* decisions (e.g. selecting goals, plans, selecting a mode) tend to be more impactful than more concrete ones. Indeed, more abstract decisions tend to influence many other more concrete decisions (e.g.

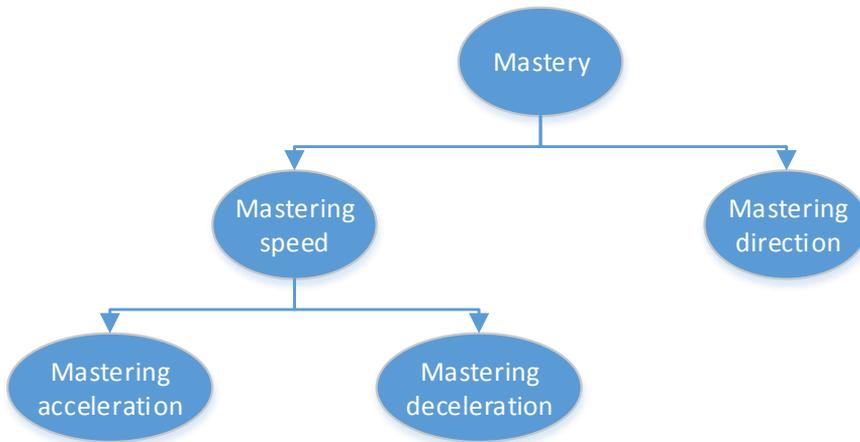


Figure 7.1 – Example of a specification tree for the value of “mastery” in the context of piloting a vehicle. As an illustration, this tree can be further specified for supporting more specific problems (e.g. “mastering direction” can be refined in “mastering turning” and “mastering altitude” for a plane)

determining the strategy “drinking tea or coffee” is generally more impactful than “drinking tea of coffee now”). In our context, we particularly focus on decisions that are expected to influence other agents or collective outcomes.

Impactful decisions provide the requirements for making value-sensitive decisions, as presented in Section 7.1.1. Indeed, impactful decisions generally lead to complex consequences. These consequences tend to be difficult to evaluate rationally, providing an opportunity for using values. Furthermore, these many consequences can be related to a rich range of concepts that can then be used for performing evaluations.

Making Abstract Values More Concrete: Specification Trees

Cultural values tend to be particularly abstract. While this abstraction is interesting for supporting a wide range of decisions, connecting them with concrete decisions can be difficult. In this section, we propose a technique inspired by the research of van der Weide (2011) for connecting abstract values to decisions.

This technique consists in connecting abstract values with more concrete values. For instance, “mastery” can be related to “mastering speed” and “mastering direction” in the context of vehicle-driving. These more concrete values can, in turn, be further specified. In some sense, a value can be expressed as a specification tree, where each node is a value and children of this node are more specific and concrete values. The satisfaction of a node should be related to the satisfaction of its sub-nodes. An illustration of that tree is proposed in Figure 7.1.

As a rule, the more children of a node are satisfied, the more the node should be satisfied. Concrete implementation of this tree depends on how values are integrated

within agent decision processes. In case of explicit evaluations, each leaf provides a concrete evaluation from the current situation and other nodes are functions that combine evaluations of their children (e.g. $\text{mastering speed} = 0.9 \times \text{mastering deceleration} + 0.1 \times \text{mastering acceleration}$). If deeply embedded within decisions, this tree can be used for raising arguments (e.g. mastering deceleration is good because it improves mastery which is good in itself).

Specification trees are a useful tool for designing value-systems. They can be used for making more concrete how value systems should influence agent decisions. These trees can be used for providing more conceptual richness and precision about the desired expression of values. In turn, this precision can be used for creating expectations.

Specification trees are useful for supporting some basic support for a wide range of decisions and more specific support for more important decisions. In the example, “mastery” is a general value that can be used in a broad range of situations. This value is further specified in the context of vehicle-driving, providing more adequate indicators for making decisions.

Creating Expectations

A last practical aspect to consider when human-like values is their affinity for supporting the creation of a richer social context. In addition of directly influencing decisions, cultural values have the core property to (be expected to) be shared with other agents. This expectation can be used in many ways, such as creating expectations about behaviors or preferences of other agents. It can be used in turn for creating a much richer social identity (i.e. norms, common plans, organization, expectations).

In particular, by being expected to be shared, values help to create expectations about other agents or about consequences of decisions (e.g. leaders can expect their orders to be completed if obedience is important). Furthermore, values strongly support the establishment of norms as highlighted by Miceli and Castelfranchi (1989). Then, these norms provide concrete instructions about how to behave with other agents, which supports efficient concrete interactions. They can also be used for creating indirect expectations about social or environmental properties. For instance, if timeliness is important, then other agents are assumed to be on time. As a result, agents can create the social assumption that schedules are reliable.

So far, no solutions for modeling and implementing the consequences of values are available from the literature. An immediate solution consists in integrating them by design. Agent designers can integrate expectations about other agents within agent decision processes, based on the value system and in considering theories of culture. Throughout this thesis, we rely on this technique and a lot of inspiration can be found in human-science theories for designing value-systems. Nevertheless, this solution requires human intervention that is not adaptable for dynamic systems.

Some directions can be considered for automatically building expectations. Game-Theoretical-like models can be expanded for including values and be used for estimating interaction outcomes (e.g. rewards can be updated according to the value systems and then be used for creating expectations about other agents). Some techniques for automatically generating norms are available, particularly for in the context of organizations (e.g. Moise+). These techniques may be expanded for including the support from values

for generating, establishing and complying with values.

Conclusion

This section investigates the core practical aspects of the integration within agent decision processes of cultural values. In particular, we consider how to represent value-systems within agents and using them for supporting decisions. In addition, we consider how to determine which decisions are adequately supported by values. Finally, we consider how value-systems, when assumed to be shared, can be more deeply integrated within agent decision processes for further improving performance.

For sure, a lot of research remains to be done to explore further how to implement the influence of values on deep cognitive aspects. This direction appears to be at the core of the design of more sociable agents, capable of better coordinating with each other.

Conclusion

This section introduces the core practical considerations for using value-systems for supporting coordination and techniques for tackling these considerations.

In conclusion, what are the practical considerations for designing value-systems for supporting coordination? First, system designers have to find adequate values for influencing a wide range of coordination-influencing decisions. Furthermore, each value should support behaviors that lead to coordination-supporting interaction patterns. Then, system designers have to put these values together and consider whether they fit with each other. Values should not overlap or conflict too much. Furthermore they should, together, cover enough coordination-influencing decisions. Next, in order to create tighter expectations about the general behavior of agents, system-designers should help the resolution of conflicting values. In order to do so, we propose to assign a relative importance to the various values of the value-system, indicating the relative importance between values. This relative importance can then be used for tackling conflicts of values. Finally, value-systems have to be integrated within agents. In order to do so, system designers should consider how they want to model their value-systems. They have to determine the type of decisions where values should be used. Together, they have to connect the value-system with decisions that agents have to make. Finally, they can consider the social context supported by values for enriching agent decision process, for instance with value-compliant expectations and behavior rules.

These steps allow designing and integrate coordination-supporting value-systems. The mechanic that makes that values support coordination has been further studied in Chapter 6. At the interaction level, human-inspired cultural value-systems indirectly support, through a direct influence on individual decisions, abstract “well-working” interactions patterns. When agents face a concrete decision situation, this abstract “well-working” pattern is implemented in a concrete interaction pattern that generally supports coordination.

7.3 Conclusion: How Can we Use Models of Human Culture as a Practical Tool to Support Flexible Coordination in Artificial Societies?

Values evaluate what seems to be “good” from a limited perspective (e.g. honesty, modesty, obedience). Values can be seen as perspective-oriented heuristics for supporting good decisions. Culture supports “cultural values” that indicate “good” socially-supporting behaviors (e.g. obedience, autonomy). In practice, these values, when they are collectively complied with, lead to better individual and collective outcomes. They prevent conflicts to occur that can damage coordination (e.g. everyone consider obedience as important, leaders give orders, subordinate obey, no conflict, better outcomes).

Value-systems are appropriate for handling multiple coordination problems. Value-systems are particularly appropriate for supporting high flexibility and robustness which is crucial for handling complex or uncertain coordination problems. They are particularly relevant for handling environments raising complex decisions, uncertainty, which can evolve with time, where unexpected things can happen. For all these problems, value-systems provide a generic support, relatively environment-independent for making decisions that maintain coherent interactions. Value-systems can be used for depicting the general social atmosphere to be found in the system, thus, more generally, for driving system performance. For all these uses, values systems *add to* existing coordination techniques. Value-systems combine relatively well with existing coordination techniques, either by providing a support when they fail or by providing a complementary perspective for supporting coordination.

On a technical side, the question consists in determining how to design such a system. In this chapter, we investigated core practical aspects to consider for designing such a system. We recognized four steps. First, designing values that support desirable interaction patterns when they are shared and are applicable for the coordination problem. Second, selecting which values to integrate such that they combine with each other with limited redundancies and conflicts. Third, ordering these values for providing guidelines for resolving conflicts. Fourth, integrating these values within agent decision processes such that they can support in practice coordination-influencing decisions.

This exploration allows looking ahead into the future. Value-systems are clearly a basis on which social reality can be constructed. This basis provides what is important driving the elaboration of more specific social mechanisms. These mechanisms can in particular be useful for concretely supporting coordination. For instance, value-systems support abstract tendencies for creating organizations (e.g. centralizing decision or not, using many rules or not). However, these tendencies can be flexibly adapted by agents for handling a specific environment (e.g. leaders can give the responsibility to assign tasks for a given problem. In another situation, the leader would manage the position of subordinates).



Conclusion

Throughout this thesis, we aimed at being capable of better coordinating agent societies when facing complex coordination problems (e.g. in complex, dynamic, uncertain environments). As preliminary research for tackling this very general goal, we investigated how complex coordination problems are solved in human societies. This preliminary research revealed that *culture* is one of the key factors for achieving coordination of individuals in human societies. Thus, in transposing human culture within artificial societies (through models), we hope to transpose the same key factors for achieving coordination in artificial societies: using culture for better coordinating agents in artificial societies.

This transposition is not immediate. Additional efforts are required for obtaining enough support from theories of human culture, using these theories for designing a model of culture and using this model properly. In order to explore it thoroughly, this transposition became the whole goal of this thesis: how to use human-like culture within artificial societies as a tool for practically supporting coordination?

In order to tackle this goal, we adopted an *exploratory* approach. This approach allowed us to cover the very broad distance separating informal descriptions from social sciences and MAS design techniques. This distance is particularly stretched by the relative lack of research relating culture with coordination. This research shortened the distance to cover by providing "shortcuts" we could have used along the way.

This exploration, from social sciences to MAS design went through several steps that we separated in two parts. The first part shows the possibility to design agent decision models that integrate the influence of culture that can be used in situations of coordination and that influence behaviors in a way that is coherent with the influence of culture in human societies. In order to do so, in Chapter 3 we first enriched further theories from social sciences with the aim of relating them more easily with MAS-oriented concepts we needed next. In particular, we focused on relating cultures with individual decisions and collective outcomes. Based on this theory, in Chapter 4 we then proposed a generic decision architecture that is culturally-sensitive and can handle decisions in the context of coordination. Finally, in Chapter 5, we showed that this model makes sense with regard to theories of culture. In order to do so, we implemented this model in a simulation context. Based on this simulation, we checked that the influence of culture on individual decisions and collective outcomes matches the influence expected to be found in human societies

as depicted in Chapter 3.

The second part shows that culture can be practically used as a tool for supporting coordination. In order to do so, we first needed to cover a theoretical gap. Indeed, theories do not explain how the known influence of culture on individual decisions lead to the known influence of culture on collective outcomes. They only describe the direct influence of culture on individual decisions and the indirect influence of culture on collective outcomes, these two relations being considered relatively independently. Since this link is crucial for understanding how culture can be used for practically supporting flexible coordination, we analyzed it in detail in Chapter 6. Finally, based on this theory and research from previous chapters, Chapter 7 finalizes the goal of this thesis. This chapter analyzes how human-like culture, more specifically value systems, can be practically used for supporting coordination.

The current chapter gathers all the insight we obtained throughout this thesis and reconsiders it. First, we can look back at the various steps we covered in this thesis, reconsider them from an integrative “bird’s view”. Based on this overview, we propose conclusive theories for explaining the relation between culture and coordination in human societies and how culture can be operationally used in artificial societies for supporting coordination. Second, we consider the progress we achieved. We reconsider the contributions we brought to the fields when aiming for our goal. We re-frame contributions with regard to the relevant research communities. Finally, we can look at the future. From the point we reached at the end of our exploration, we can see better what can be done next. In particular, we can indicate how the path that we explored during our exploratory approach can be further cleared, but also what are the targets that are in sight for the future.

This chapter is structured as follows. Section 8.1 introduces the core insight about the relationship between culture and coordination, that we obtained throughout this thesis. Section 8.2 introduces our contribution to the fields. Section 8.3 presents future research that can be foreseen from the current thesis.

8.1 Answering Research Questions

This conclusive chapter provides a good opportunity for reconsidering existing theories with the addition of the progress made throughout our thesis. To this end, we answer the research questions we formulated in Section 1.3. Each following subsection addresses one research question.

8.1.1 Answering Research Question 1: *How Does Culture Influence Coordination in Human Societies?*

This question aims at understanding better how culture influences coordination from the perspective of human societies. In particular, the answer to this question should provide a conceptual basis detailing core aspects and core expectable properties of the influence of culture on coordination. This basis is to be used later for building models.

This question raises three more specific sub-questions. These sub-questions raise the core concerns for designing MASs: influence of culture on individual decisions, the influence of culture on collective outcomes and the link from individual decisions to collective

outcomes.

The current section summarizes and reconsiders these three subquestions based on existing works and expanded with discoveries made throughout this thesis. Each subsection answers one sub-question. Then, a final subsection proposes a further expansion of our theory in the light of our latest insight obtained through generative social science.

Answering Research Question 1.a: *How Does Culture Influence Individual Decisions Regarding Coordination in Human Societies?*

Similar mechanisms of the influence of culture on individual decisions are at play both in the general case and for more specific decisions. Thus, in the current section we first consider the influence of culture on individual decisions in general and then we specifically focus on coordination-related manifestations.

According to Hofstede et al. (2010a), the main theory we rely on, the influence of culture on individual decisions can be explained by using two cognitive aspects: *practices* and *values*. Practices can be seen as collectively-shared behavioral and symbolic standards. They are composed amongst others of symbols, heroes and rituals.

Values strongly influence what individuals can consider as “good” (e.g. obedience, autonomy, security, creativity, assertiveness, modesty, honesty). Values are described by Miceli and Castelfranchi (1989) as cognitive mechanisms that evaluate the “good qualities” of something (e.g. all other things being equal, honest answers are better than dishonest ones). When making decisions, values can be seen as heuristics that are irrational, but generally accurate and quick-to-evaluate. They are generally used for cutting short decisions that raise difficulties for rational decision mechanisms (e.g. consequences that are hard to predict or to compare).

Values are also crucial in the context of interactions. First, because decisions in situation of interaction raise uncertainty, which pushes more easily rational decision processes to their limits. Second, because values support the occurrence of consistent qualities in one’s behaviors. This consistency allows creating expectations about others (e.g. believe someone who gives a lot of importance to honesty). This aspect is particularly important in the context of culture, where values are expected to be shared, allowing to create strong *a priori* expectations about other’s behaviors. Third, because values can be concretely instantiated in many ways. They can be used as an abstract framework for determining the qualities that someone is interested in. This framework can then be instantiated in many ways for interacting (e.g. making the right offer, finding common interests). Miceli and Castelfranchi (1989) are particularly insisting on the instantiation of values into norms, which can then support concrete interactions (e.g. I apply a norm forcing to report information to leaders because this norm supports obedience and obedience is “good”). The support provided by values is even more interesting for supporting interactions when values are (culturally) shared.

Culture strongly influences the set of values adopted by (a group of) individuals and the relative importance given to values, which we refer to as the *value system*. This relative importance matters when making decisions for which various values support different alternatives (e.g. “Shall I cover up a fault of my boss?” Available alternatives are linked to either honesty or obedience. The relative importance between the two is at play).

In considering values and practices together, both influences are expected to be shared

with other individuals with the same culture. Thus, both can be used for creating expectations about others. In a given culture, values and practices tend to match each other (e.g. high importance to the value of obedience and discipline generally matches to means for differentiating status of power in natural language). Values are recognized to be at the core of culture. Values are harder to change than practices. Values influence a broader range of decisions than practices.

Through HCDs, Hofstede et al. (2010a) empirically show that culture tends to influence the expression of specific types of behaviors and which can be related to some specific *decision aspects* (e.g. importance of statuses, tolerance of uncertainty, acceptability of failures). In terms of values and practices, HCDs indicate abstract types of value systems and practices supported by a given culture (e.g. societies giving high importance to status are more likely to promote values of obedience and discipline). In this direction, Chapter 3 enriches the descriptions from Hofstede et al. (2010a) by relating the type of values that can be expected to be found and supported by culture given their scores on HCDs. This relation allows connecting values with concrete decision aspects and behaviors sensitive to HCDs and vice-versa.

HCDs and the sets of decision aspects they highlight can be used for considering how culture influences decisions in the *context of coordination*. In this context, D'Iribarne (1989); Hofstede et al. (2010a) proposed numerous concrete examples of decision situations where culturally-sensitive aspects are at play and influence coordination outcomes (e.g. in a high PDI-CD culture, subordinates prefer to rely on common leaders for solving conflicts). Chapter 3 used these examples for proposing more generic culturally-sensitive decision rules (e.g. in high PDI-CD culture, opinion of leaders are given more importance than those of subordinates. Thus, opinions of leaders are appropriate for settling conflicts between subordinates). These general rules can then be used for explaining (and designing) more easily decisions made by individuals in the context of coordination.

In addition to these generic aspects, Chapter 6 introduced additionally a new aspect of the influence of culture: *cultural expressiveness*. The coordination context (environment, coordination structures) influences how much and which parts of culture are expressed. Indeed, culture influences a limited set of decision aspects. These aspects may match (or not) with the set of decision aspects which are relevant for the given coordination problem (e.g. cultural sensitivity to uncertainty is applicable only when the environment is uncertain). We identify multiple causes that may influence the expressiveness of culture. First, the decider needs to have enough *freedom*: decisions should allow multiple answers, or alternatives. Second, the various alternatives for a decision should be *conceptually related* with the cultural background of the decider. Third, culture should help *discriminating the various alternatives* of the decision to be made (e.g. honesty hardly influences the decision "tell the truth" vs. "write down the truth"). Fourth, in order to be sensitive to culture, decisions should involve a *rational dilemma*. In rational dilemma, rational deliberation can only provide limited support for determining the best alternative for making the decision (e.g. alternatives are hard to rationally order, finding the best outcome requires too high deliberation effort with regard to available time for making the decision). In such decision situation, it makes sense to rely on (irrational) values for providing relatively helpful support for making decisions.

Finally, individuals can possess multiple cultures which are expressed depending on the context (e.g. at home, at work, at sports). While other aspects of culture had to be

tackled first in this thesis, we think that this aspect opens a crucial gate for further conceptualizing and implementing the influence of culture in situations of coordination. This aspect is used later in this chapter for highlighting new directions.

Answering Research Question 1.b: *How Does Culture (Indirectly) Influence Collective Outcomes in Situation of Coordination in Human Societies?*

Hofstede et al. (2010a) highlight numerous relationships between culture and collective outcomes. This influence is indirect, since culture is assumed to influence individuals only indirectly, through values and practices. While the causes of this link are explained in the following section, we introduce here some core aspects of the indirect influence of culture on collective outcomes.

Former theories such as Hofstede et al. (2010a) strongly suggest that sharing culture leads to better coordination outcomes (e.g. better mutual understanding, avoiding failures). More specifically, a whole chapter of Hofstede et al. (2010a) is dedicated to illustrating the negative consequences of cultural mismatches (e.g. cultural clashes in the form of mis-expectations and misinterpretation leading to miscommunication and distrust) and how to manage cross-cultural encounters. This apparent benefit from using culture was actually the starting point of this whole thesis.

Nevertheless, the statement “culture support coordination” was at the state of intuition, while this point is crucial for our thesis. Through our studies performed in Chapter 6, we can now reconsider more sharply this statement. We found out that, in accordance with theories, sharing the same culture does support coordination *in most of the cases* (e.g. by avoiding misunderstandings). Nevertheless, sharing the same culture *does not always* improve coordination. First, because misunderstandings raised by cultural mismatches can “by accident” lead to better outcomes, even though these “good accidents” can be expected to be exceptional. Second, because some cultural features can be inappropriate for handling some coordination problems (e.g. obedience supports centralization that is inappropriate in environments where communicating is too expensive). In other words: coherent interactions supported by culture lead generally but not necessarily to better performance.

Existing theories, such as Hofstede et al. (2010a), implicitly suggest that culture influences coordination relatively independently of the coordination context. These theories generally describe the influence of culture on collective outcomes without introducing a very specific coordination context (e.g. high MAS-CD tends to support higher efficiency but lower adaptiveness to customer’s needs). This description lets believe that the influence of culture on coordination may be independent on the coordination context.

We further study this independence relationship in Chapter 6. In agreement with existing theories, we found out that culture influences similar types of interaction patterns and thus similar collective outcomes, relatively independently of given coordination settings (e.g. obedience supports centralization, which supports in turn good performance for simple problems and lower robustness). Nevertheless, we moderate this independence by adding a variable: cultural expressiveness. Individuals need to be capable of expressing the cultural influences which support the interaction pattern which support the collective outcome. For instance, the influence of obedience on collective outcomes can be expected to be more influential in pyramidal organizations than in flatter ones.

Answering Research Question 1.c: *How Does the Influence of Culture on Individual Decisions Lead to the Influence of Culture on Collective Outcomes in Human Societies?*

Being capable of determining how collective outcomes result from individual influences is crucial for building representative MASs. Unfortunately, former theories did not provide rich explanations about this phenomenon.

This question is further investigated in Chapter 6. In this chapter, we explain that culture, through values, supports the occurrence of *abstract*, *coherent* and *recurrent interaction patterns* (e.g. obedience supports centralization of decision and information around leaders). By coherent we mean that individuals are culturally-driven towards achieving these interaction patterns. If they reason about the interaction they conduct, individuals may even explicitly try to pursue these interaction patterns. Based on this, they can create strong expectations that these interaction patterns will occur and thus have expectations about behaviors of other agents. In general, interaction patterns supported by culture are Nash dilemma: individuals who deviate from them can expect lower individual (and collective) results. By abstract and recurrent, we mean that culture supports an abstract form of interaction (e.g. centralizing decision power to leaders). This abstract form of interaction can then concretely instantiated for specific interaction-related decision situations (e.g. let leaders set assignments in a contract-based interaction). To that extent, all these concrete instantiations tend to repeat the same recurrent abstract coordination mechanism.

The influence of culture on individual decisions supports the occurrence of interaction patterns supporting in turn the occurrence of known collective outcomes. The property of coherence supports the prevention of cultural clashes, which happens when culture is not shared (mis-expectations and misinterpretations, e.g.: both leaders and subordinates expect the other to manage the workload of subordinates). In turn, preventing cultural clashes leads to generally, but not always better, collective outcomes. Furthermore, by supporting the occurrence of similar abstract interaction patterns in a wide range of situations, culture supports in turn a general “direction” that leads to the known influences of culture on collective outcomes.

Latest Insight: From Culture to Cultures?

Throughout this thesis, we considered the relation between culture and coordination from a very pragmatic approach, trying to remain as close as possible to available theories. In this last section, we propose to look ahead for a seemingly coherent yet to-be-validated more complete explanation. This explanation is based on all the research performed throughout this thesis, particularly with insight raised by practical concerns for using culture for supporting coordination.

In this thesis, we focused on the most generic forms of culture. These forms provide a basis for abstractly supporting decisions and are applicable in the widest range of decision situations (e.g. very abstract values, practices for generic interactions: how to greet, how to eat). This culture can, as an image, be related to the “national culture” from Hofstede et al. (2010a).

While providing a generic support for most of the decisions, abstract cultural aspects are neither adequate nor practical for supporting decisions in more specific (coordina-

tion) contexts (e.g. decisions on the production line, at the accountancy office). In these contexts, individuals appear to behave according to a more specific culture (e.g. with more specific values such as “reporting problems to leaders” and more specific practices such as specific protocols for reporting information). This more specific culture provides more appropriate and specific elements for supporting decisions in the given context. As a matter of comparison, this culture can be seen as the “organizational culture” depicted by Hofstede et al. (2010a).

Since culture is context-dependent, multiple contexts can support the creation of multiple cultures (e.g. “report problems” can be a very specific value in the context of the production line while “report results” is another in the context of the accountancy). Each context favors the expression of related cultures. Individuals can be influenced by multiple cultures at a time when multiple contexts are interleaved (e.g. an individual with a culture of “the good example” at home and another of assertiveness at work, having to make decisions with his or her child present at the office).

The more abstract cultures that we studied in this thesis influence the constitution and the expression of those more specific cultures. Indeed, when individuals are confronted to a new context, they are more likely to adopt context-specific values and practices which match their own more abstract value system. For instance, if “obedience” is important, individuals are more likely to adopt “report information to leaders”¹ as a value. Likewise, more specific culture-compliant practices can be elaborated. For instance, they are more likely to accept protocols for providing information to leaders (driven by the value of obedience), possibly based on more generic practices (e.g. an instance of the generic ritual of “sending a report”).

In conclusion, we do not think that there is only “one” culture. Instead, we think that individuals possess multiple cultures from and expressed in the various contexts they face, possibly with various scales of abstraction and genericity. The more abstract cultures, – that we studied throughout this thesis –, provide abstract and generic directions for agents. Based on this influence and on specificities of the interaction context they face, individuals can elaborate more specific cultures for supporting context-specific interactions. While abstract cultures do not directly determine how to conduct concrete interactions, they indirectly support the elaboration of more specific cultures that then support these interactions.

8.1.2 Answering Research Question 2: *How to Model the Influence of Culture on Coordination?*

Research Question 1 shows how culture relates to coordination in human societies, particularly focusing on the support provided by culture to coordination. Next, the challenge consists in replicating similar (desirable) properties in artificial societies. In order to do so, we rely on modeling that we express through Research Question 2.

In order to answer this question, we investigate first that are the core aspects of culture to be modeled. Then, we consider how to technically model these aspects.

¹See specification trees in Section 7.2.4 for an illustration of such a specification.

Answering Research Question 2.a: *What Are the Key Mechanisms of the Influence of Culture on Coordination?*

Culture is a broad and fuzzy topic that is hard to delimit and thus to model. In order to build adequate models of culture, we should delimit the key aspects of culture which are at play when culture supports coordination. In particular, these aspects should be useful for supporting decisions in the context of complex environments.

Throughout this thesis, we recognized some key aspects to be considered for modeling culture in our context. Nevertheless, in spite of having considered many aspects, we cannot claim to provide a complete picture. Indeed, the number of approaches for investigating the topic of culture is huge and most of these approaches generally remain informal, making hard to evaluate what their models may provide. Furthermore, there may be further aspects for considering culture that will be discovered later. While we do not guarantee to be complete, we still have found several key aspects of the influence of culture in the context of coordination.

Former research, introduced in Section 2.2.4 provides some models of culture influencing individual decisions. These models are mostly based on HCDs: culture is represented by cultural dimension scores that influence very specific agent decisions (e.g. PDI-CD influences the decision to defect when interacting with an organization with higher status). This approach is behavior-oriented: these models remain as close as possible to empirically-validated cultural behaviors. While providing insight about important aspects of culture, cultural-dimension-based approaches are inadequate for our purpose. Indeed, they are not appropriate for handling the inherent complexity raised by coordination problems. These problems require richer “cognitive” decision processes.

In this thesis, we expanded this former research by considering *value systems* as a key aspect of the influence of culture in human societies in situation of coordination. The answer to Research Question 1 provides a rich theoretical background for considering the importance of value systems. Value systems are shown to be at the core of human culture (literally, in the onion diagram from Figure 2.2). In addition, value systems are shown to be particularly appropriate for supporting complex decisions. For the sake of checking that value-systems are applicable for replicating core aspects of human culture in the situation of coordination we modeled multiple simulations. Chapter 5 and Chapter 6 show that value-based models of culture are sufficient for reproducing individual and collective behaviors that match those observed in human societies.

We also considered practices, which are the second component of culture presented by Hofstede et al. (2010a). Nevertheless, practices appeared to be less relevant for supporting coordination than value systems. First, the support provided by values is more restricted (need for specific situations). Second, existing coordination techniques already cover many aspects of practices (e.g. rituals, symbols). Thus, using practices would be redundant with these techniques, failing to offer additional benefits.

Answering Research Question 2.b: *How Can We Build Practical Models of the Key Mechanisms of the Influence of Culture on Coordination?*

Once the key aspects for modeling culture have been identified, the next question consists in modeling them.

In considering available work, the best computational value-sensitive decision model we know of is provided by van der Weide (2011). This model supports decision-making that is in line with underlying values. Nevertheless, this model focuses on limited one-shot and self-oriented decisions. This model is not suitable for our setting, which involves making series of decisions that require to take others into consideration.

In this direction Chapter 4 expands the model from van der Weide (2011) by integrating value systems within the BDI architecture. The BDI architecture is particularly adequate for supporting practical reasoning, which is appropriate for making decisions in the context of coordination. Then, values can complement the BDI framework for making decision when multiple rationally-similar alternatives are available. Values can thus support the BDI framework when complex decisions are to be made. In this chapter, we also suggest multiple decision aspects that can be supported by values.

In addition to this architecture, we introduce in Chapter 7 a series of techniques for implementing it. In particular, we reintroduce a modeling technique from van der Weide (2011), the design of specification trees. This technique can be used for connecting abstract values into more concrete values that can be easier evaluated concretely.

As a closing remark, we want to highlight that in spite of the progress made in this thesis, modeling value-systems for making complex decisions is still stammering. So far, available models support “value-compliant decisions”, decisions that are in line with values. Nevertheless, values can be used as a much richer tool for making decisions, as highlighted by the many interesting decision properties of values introduced by Miceli and Castelfranchi (1989) (e.g. automatic generation of values; decision heuristics; normative support).

8.1.3 Answering Research Question 3: *How Can We Use Models of Human Culture as a Practical Tool to Support Coordination in Artificial Societies?*

In order to further answer this question, we need to take into consideration the context in which this question has been raised. We want to consider how *system designers* can use culture (more specifically, value systems) for supporting flexible coordination. System designers have limited resources (e.g. time, money) for coordinating agents. Thus, the support provided by culture should, at least for some decision situations, lead to better results than the support offered by other coordination techniques.

In this thesis, we tackled this initial question through the investigation of two crucial sub-questions: when are value systems appropriate for supporting coordination and what are the practical considerations for designing culture for supporting flexible coordination. To our knowledge, none of these questions has been investigated by previous research.

Answering Research Question 3.a: *When are Value Systems Appropriate for Supporting Coordination?*

Our investigations conducted in Chapter 7 indicated that value systems are appropriate for supporting coordination in *complex and dynamic environments*. This complexity can be characterized by multiple aspects such as uncertainty, complex evolutions, many different types of decisions, conceptual complexity, computational complexity, high level of

interactivity between agents, high number of agents and interactions with humans. As highlighted by former research questions, in these environments, rational reasoning is limited by this inherent complexity. This rational reasoning is complemented values that providing heuristics for supporting (rather) good decisions.

In these environments, traditional coordination mechanisms (e.g. norms, protocols) are limited by inherent complexity. Indeed, these mechanisms are meant to coordinate agents “externally” or *a priori*. Coordination is exclusively done by designers. However, in complex environments *a priori* coordination requires to take into consideration far too many aspects and possibilities. Thus, design complexity increases too fast, leading to impracticality.

Some more flexible coordination techniques (e.g. abstract organizations, cooperation) are more adequate for handling this kind of environments. These coordination mechanisms provide the basic means that agents require for coordinating and let agents manage coordination by themselves. In this case, agents can coordinate on-the-fly, when facing a concrete instance of the environment.

The support provided by values complement the one provided by those techniques. In particular, values have some properties that other coordination mechanism do not provide or only partly, such as consistency, mutual expectations, collective agreement about desirable outcomes and abstract coherent interaction patterns can be made concrete.

In this conclusive chapter, we can relate this method with regard to the key aspects for considering coordination raised in Section 2.1.3. Our method is adequate for handling uncertain and complex environments. It can be used as an explicit coordination constraints and can be provided for supporting coordination in open systems. Our method is strongly coupled with goal autonomy and dynamic social configuration. Our method tends to limit emergence and discards locus of control, even if it can be added as a value. Finally, our method supports mean efficiency (costs for internally coordinating), high flexibility and variable robustness.

Answering Research Question 3.b: *What Are the Practical Considerations for Designing a Culture to Support Flexible Coordination?*

Our investigations conducted in Chapter 7 describe several concrete technical solutions for practically managing using culture for supporting coordination.

We introduce four crucial steps to be taken into consideration when designing value systems for supporting coordination. These steps are: designing values, selecting which values to integrate, determining how to combine these values into value systems and finally integrating value systems within agent decision processes.

In addition, we introduce crucial aspects to take into consideration when designing value-systems for supporting coordination. Value systems should be computationally efficient, cheap enough to design, be driven towards supporting coordination and be conceptually in line with human inspirations.

We propose concrete techniques for implementing value systems for each steps while taking the various aspects into consideration.

8.2 Contribution

Throughout our exploration towards tackling the goal set by this thesis, we contributed to multiple scientific communities. Each following subsection introduces the contributions we offered for each community we contributed to.

8.2.1 Contribution to Social Sciences

Our exploration started from theories from social sciences depicting the influence of culture on humans and human societies. Throughout this thesis, we contributed to these theories by enriching them in order to get better understanding of the influence of culture on coordination.

Our first contribution to this community is proposed in Chapter 3. This chapter expands existing theories by proposing a conceptual generalization of available theories. These latter theories were only proposing limited examples for relating the influence of culture on individual decisions and on collective outcomes. We needed more generic influences for proposing more exhaustive individual models of culture and creating expectations about collective outcomes.

Our second contribution is also proposed in Chapter 3. This chapter relates conceptually and empirically the core values to be expected to be promoted by culture depending on HCD cultural scores. While former research has drawn relationships between the two, no exhaustive list was available.

Our third contribution is proposed in Chapter 6. Current theories describe how culture influences individual decisions and how culture influences collective outcomes. Nevertheless, these theories do not explain how the influence of culture on individuals decisions leads to the influence of culture on collective outcomes. Chapter 6 uses generative social sciences Epstein (1999) for suggesting an explanation to this link by considering an additional perspective: interactions. We explain that each culture seems to support the occurrence of some coherent, recurrent and abstract interaction pattern. This culturally-supported interaction pattern appears to result from the influence of culture on individual decisions, either implicitly or explicitly. In considering collective outcomes, we can relate more easily the culturally-driven interaction with culturally-driven collective outcomes (e.g. good efficiency for simple problems but bad robustness supported by obedience can be explained in terms of the centralization supported by obedience). As a matter of illustration, several examples of the links between values, individual decisions, interactions, and collective outcomes are presented in Chapter 7.

8.2.2 Contribution to the Agent Design Community

Our thesis required designing an agent decision architecture that is both value-sensitive and appropriate for making decisions in coordination situations. These two aspects are part of the core topics investigated by the agent design community.

Chapter 4 proposed an abstract decision architecture that integrates both BDI and values. This architecture is particularly appropriate for making complex (e.g. tight interactions, complex environment) and goal-directed decisions. The traditional BDI formalism

is particularly relevant for performing practical reasoning and for determining the adequate decisions to tracking a goal. Nevertheless, this formalism relies on rational processes. This rationality raises limitations when decisions become too complex (e.g. computational complexity for determining the best outcomes). This limitation can be tackled by values. Indeed, values are irrational yet principled and sound decision heuristics. These heuristics can cut short endless deliberations by raising easy-to-evaluate indicators of a “good” solution. To our knowledge, this architecture is relatively new to the community. van der Weide (2011) proposes a model for supporting practical reasoning with values, mostly focused on argumentation. Our contribution consists in explicitly connecting this model with the BDI architecture, for achieving better practical reasoning.

Chapter 7 enriches existing theories for concretely designing value-sensitive decision processes. So far, designing value-sensitive decision processes remains a relatively open problem. Miceli and Castelfranchi (1989) propose the core theoretical foundations of cognitive aspects involved in values and van der Weide (2011) proposes solutions for modelling them, but this task remains far from complete. Based on this previous research, we suggest several (pragmatic) solutions for implementing such a decision model. We particularly consider the use of values as provided by human culture, generally driven towards interaction-related decisions.

8.2.3 Contributions to the Simulation Community

Our thesis required to model and study credible human-like culture. In order to check the credibility of our models and for analyzing this culture, we designed multiple simulation models. These simulation models adds to the simulation community.

Chapter 4 proposed an original model of the influence of culture on values. This model enriches the set of available cultural models which are used in the simulation community. Available models focus on reproducing simple behaviors, generally using HCDs as an inspiration. However, this simplicity limits the depth of decisions that can be influenced by culture: agents are limited to performing simple decisions. On the contrary, values allow supporting much more complex decisions because they are more deeply cognitively embedded. Thus, models of culture based on value systems can be used for agents capable of making decisions that are more complex. Thus, our value-based cultural model offers the possibility of using culturally-sensitive agents in a wider range of simulations (e.g. rich agents for serious gaming, using values for making and explaining decisions as suggested by van der Weide (2011)).

Chapter 5 and Chapter 6 propose two simulation models that are validated against expectations from social sciences. These simulations provide concrete examples for the social simulation community. Furthermore, these simulations relate two topics that have not been extensively considered together, yet are useful for making simulations: culture and coordination.

8.2.4 Contribution to the MAS Design Community

The end-goal of our thesis consists in managing better to coordinate artificial agents. This goal is also one of the core goals of the MAS design community.

We contributed to this community in Chapter 7. We proposed to use value systems for coordinating agents. To our knowledge, this technique has never been investigated before. This technique appropriately helps to tackle one of the main challenges of the coordinating agents: achieving coordination in complex environments (e.g. uncertainty, inherent complexity, dynamism, open environment, interaction with humans). Our work provides a new means towards being capable of keeping coherent interactions in a very wide range of environments, paving the way towards agent societies capable of higher autonomy.

8.3 Recommendations for Future Research

The exploration conducted in this thesis touched multiple topics: social science theories, designing and simulating culture and culturally-sensitive agent-based decision models and using culture as a tool for supporting coordination. Our exploration reveals for each topic additional interesting research tracks that can be explored next.

Three perspectives can be considered for expanding each topic. First, additional research can further strengthen or improve our exploration (e.g. more complete validation, techniques for using our models in a better way). Second, our exploration highlights possible crossroads that can be explored next. While indicating interesting research directions, these crossroads had to be left out because they do not lead to achieving the goal of this thesis. Third, future work can continue our exploration, by proposing further tracks that support and expand our theory.

Each following section focuses on a topic that we went through. Each topic is then split in three subsections: from exploring to paving the way, crossroads and further exploration.

8.3.1 Social Science Theories

Further research can enrich our work in the topic of social science theories. Furthermore, our research highlights additional research directions. Finally, future research in social science theories can help to expanding further the content of this thesis.

From Exploring to Paving the Way

Additional research can strengthen social science theories proposed within this thesis (in Chapter 3 and Chapter 6). The most direct possible improvement consists in further empirically validating our theories, for instance through polls or case studies. Indeed, our theories are based on expansions of other theories and investigations through simulations. While our theories match with regard to existing research, some parts of our theory could not be compared against detailed empirical findings, because simply there are no such empirical data. By performing these empirical studies, we will obtain additional means for validating and enriching our theories and thus subsequent models. As a specific point to investigate, we recommend in particular investigating in details cross-relationships between the influence of culture on individual decisions, interaction patterns and collective outcomes, from a holistic standpoint. For instance, several detailed case study *à la* D'Iribarne (1989) that record and compile detailed empirical data at the individual, interaction and collective level, formally and informally, including culturally-sensitive aspects

such as values would give a solid ground for relating these aspects altogether and building more accurate models.

Crossroads

Chapter 6 introduces a theory for connecting known relations about culture and coordination, from individual to collective level. This theory introduces two fresh directions from the topic of social sciences: this theory proposes a bottom-up explanation and this theory relates culture and interactions. These new directions provide new research perspectives for understanding culture better.

First, we show that the influence of culture can be explained bottom-up, strongly relating together individual and collective levels. This perspective is particularly adequate for proposing integrative theories. These integrative theories aim at capturing underlying mechanisms of culture. Then, these mechanisms can be used for explaining empirical observations about culture and better supporting predictions that can be made. This perspective is crucial for enriching current theories of culture. As highlighted in Chapter 2, current theories of culture focus on limited (indirect) correlations (e.g. this collective phenomena happens in this culture). Focusing on correlations leads to highly fragmented bits of knowledge, making hard to investigate, understand and explain cultural phenomena. Conversely, integrative perspectives offer new means for understanding the influence of culture, which deserve to be further explored. As non-negligible positive side effect, understanding the underlying mechanisms of culture is directly applicable for proposing models of culture that are more accurate and realistic. Back to theory, these models then be used for building simulations and providing cheap yet comprehensive validation of underlying theories, as we did in Chapter 6.

Second, we highlighted the influence of culture on interactions. The perspective of interactions is relatively unexplored and offers interesting research directions. In particular, Our theory can be expanded for better understanding influence of culture in general, not only in the situations of coordination.

Further Exploration

Future research from social sciences can be helpful for better using culture as a tool for supporting coordination in artificial societies. In this setting, social science theories are useful for forming an idea of the relations between culture, individual decisions, interactions and collective outcomes in the situation of coordination.

Future research can expand and refine the theories proposed in this thesis. While being sufficient for having clear-enough ideas for achieving our goal, our theories remain raw. Refining them can provide clearer insight about how culture relate to coordination. Then, this insight can help us for better determining how human-like culture can support coordination in artificial societies and what to include within models of culture.

Some core directions can be given as shorter-term goals. First, cognitive mechanisms involved in the expression of culture remain relatively unexplored. Values and practices provide basic explanations but more is expected to be found further on this track. This research is particularly relevant for credibly modeling culture and designing culturally-sensitive agents capable of handling a wide range of decision situations. Second, further research can further relate culture to interactions, particularly through the influence of

culture on individual decisions. These links will provide better expectations about concrete manifestations of culture to be observed in artificial societies. They can help, for instance for designing adequate culturally-coherent coordination mechanisms.

8.3.2 Modeling and Simulating Culture

Models and simulations proposed in this thesis open new perspective for future research. This research can be useful for further improving current models, opening other research directions or introduce new solutions for using culture for supporting coordination.

From Exploring to Paving the Way

When designing the models and simulations proposed in this thesis, we faced two major limitations. Solving these limitations can help improving our models and simulations.

Lack of Data: More adequate empirical data will help both the design and the validation of future models. One of our major difficulties when designing models was the lack of coherent databases. Most of the available data provide fragmented bits of insight, but they fail to offer coherent wholes. When making simulation models, we suffered from the lack of “concrete examples” to hang on and use as a target for producing a model. The only exception we had was the work from D’Iribarne (1989), which proposes relatively deep comparative case studies of three similar organizations from different countries. However, more effort is required for obtaining more formal and comprehensive data.

We acknowledge the difficulty of gathering cross-cultural coordination-related data. Nevertheless, more research in the direction of D’Iribarne (1989) opens the gate for building more accurate and realistic models of culture. Such a study would propose in-depth rich information about individual values, individual decisions, interactions and collective outcomes. We hope that the increasing importance of embedded information technologies will enable the collection of such data. This aspect is crucial for improving the design of models and simulations.

Lack of Concrete Means for Implementing Value-Sensitive Decisions: Our model of culture relies on the influence of value-systems on individual decisions. Nevertheless, so far, computational solutions for modeling value-sensitive decision processes remain limited. In this thesis, we worked around this problem by proposing an abstract decision model and implementing the influence of culture “by hand”, through value-driven hand-written decision rules. Nevertheless, being able of better designing adaptive value-sensitive decision models can significantly expand the range of applications that can be handled by our value-based culture approach.

In considering possible tracks for future research in this direction, some limited computational solutions are available, such as the work of van der Weide (2011) and the insight we provide in Section 7.2.4. More conceptually, Miceli and Castelfranchi (1989) introduce a rich definition of values and how they influence other cognitive mechanisms. For instance, values are used by individuals for creating heuristics about “good” aspects of decisions; individuals can generate values on their own, through extrapolation from previous experiences.

Crossroads

We offer a new model for integrating a human-like influence of culture for performing complex decisions. This model can be applied in a wide range of simulation settings, where agents need to be sensitive to culture and have to make relatively complex decisions. Thus, our model opens the way to a new range of simulations of culture. A particularly relevant target for value-sensitive decision models is serious games as proposed by van der Weide (2011). This setting can require agents capable of performing complex (yet not necessarily rational) decisions and of explaining their decisions in terms that humans can understand.

Further Exploration

Additional directions can be explored for designing new models of culture that can be useful for supporting coordination.

The most immediate direction arises from the theory of Hofstede et al. (2010a). This theory introduces an additional cognitive aspect of culture: practices. In the context of coordination, practices are particularly useful since they provide concrete means for integrating the influence of culture on decisions and for creating tight expectations about behaviors of other agents. Some design solutions conceptually match design practices. For instance, rituals can be represented by protocols and symbols by ontologies. Modelling heroes is still under question and will require future exploration (e.g. a “hero” agent, visible by other agents, would be highlighted as an example to be learned from).

Additional extensions can be proposed (e.g. learning, integrating multiple cultures). They are further related to theories for using culture for improving coordination and thus are presented in the following section.

8.3.3 Using Culture for Supporting Coordination

Finally, our target, using culture for improving coordination, is not a definite end-goal. This theory can be further explored and opens the way for future models and techniques.

From Exploring to Paving the Way

The most straightforward next step for strengthening the theoretical framework presented in Chapter 7 consists of concretely applying this theory. In other words, building a working applications where culture is used for supporting coordination in an artificial society.

Ubiquitous computing is an interesting topic for introducing culturally-sensitive agents. Ubiquitous computing offers a rich ground for concrete application that features both human-agent interactions and tight interactions. Values are particularly human-friendly and thus appropriate for achieving fruitful human-agent interactions. First, because values provide a handy communication interface with humans (e.g. arguing with values from van der Weide (2011), using values for justifying norms from Miceli and Castelfranchi (1989)), which is one of the core challenges of human-agent interactions. Second, values provide a useful framework for determining what humans consider as important. This information can then be directly tackled by agents. Third, values support informal

abstract interaction structures that are particularly appropriate in human-agent interactions. These informal structures provide high-level indications about how to conduct interactions. These indications are generally used by humans for interacting, against for instance detailed protocols. Furthermore humans can use these indications for creating expectations about how these agents interact with each other (e.g. what means being the owner of the system? Being the “leader” giving instructions to the system or the “customer”, hoping to be well served by the system?). Additionally, for systems dedicated to single users or users with similar values, these values can be taken into consideration when designing the system. This possibility enables value-driven design, for better fitting user desires and expectations, in the direction of the research from Friedman et al. (2006). Finally, values are useful for designing agents capable of adapting to a wide range of situations. To that extent, humans can engage more easily in a wide range of interactions with agents composing this artificial society. This wider range of possible interactions improves the usefulness of the system as a whole.

Ubiquitous computing also features rich interactions between agents while keeping a relatively low need for individual adaptiveness to the environment (e.g. limiting an agent to managing a fridge is simpler in terms of self-oriented environment-management than exploring Mars while still raising complex interactions with other agents). This high level of interaction requires agents to adapt to many other agents, which is an appropriate setting for using value systems.

Further Exploration

Additional directions can be explored for better managing to use culture for supporting coordination. We particularly foresee two core directions.

Multiple Cultures: As theorized in Section 8.1.1, individuals seem to possess multiple context-dependent cultures (e.g. a culture of assertiveness as work, another one of kindness at home). These cultures are created and expressed depending on the decision context. Cultures have various levels of abstraction and genericity. More specific cultures tend to subsume generic cultures in the adequate context (e.g. at the workplace, the influence of work-related culture is more influential than the more generic national culture). Still, more specific cultures remain sensitive to more generic cultures (e.g. if obedience is important in the national culture then individuals are more likely to promote obedience-related values and practices in their culture at work). Finally, the influence of context-specific cultures tend to be more specific than more generic cultures (e.g. the value “reporting information” from the culture at work is more specific and applicable to the work context than the value “obedience” from national culture). Thus, context-specific cultures support the creation of tighter mutual expectations between individuals.

The possibility for integrating multiple context-dependent cultures improves the possibilities for supporting coordination. Generic cultures would support a broad range of decision situations, providing a backup for exceptional situations. More context-specific cultures, based on stronger expectations about the situation, would support tighter interactions and mutual expectations. Thus, context-specific cultures are appropriate for improving efficiency for more expectable decision situations.

In a longer-term, we can consider the possibility for (semi-)automatic generation of context-specific cultures. Abstract cultures provide a generic interaction framework that

can be used by agents for handling basic interactions. When agents consider that they are going to interact with other agents for a specific context (e.g. recurrent interactions), they can further refine this generic interaction framework for locally optimizing their interactions with others. In some sense, they would enrich their social environment. By this means, agent societies would manage both to coordinate in a wide range of environments and to achieve higher performance for more context-specific recurrent interactions. This mechanism seems to reflect what happens in human societies and supports the occurrence of local sub-societies. The integration of such a mechanism may open the gate for a higher level of autonomy for agent societies.

Evolving Culture

In this thesis, we assumed for the sake of simplicity that culture is static. Nevertheless, human cultures evolve, over time depending on numerous other factors (e.g. interactions favored by society and the environment) as explained in Hofstede et al. (2010a). Some theories such as the one of Beck and Cowan (1996) further investigate the evolution of culture in the context of coordination.

Beck and Cowan (1996) show the mutual-dependence relationship between culture, societies and environmental conditions. They explain that the environment and social conditions (e.g. social size, social structure) support the occurrence of some type of generic abstract cultural patterns. These cultural patterns support in turn the occurrence of abstract interaction patterns, influencing in turn environmental conditions and social structure. For instance, groups from few hundreds to few thousands of individuals tend to organize in chiefdoms. Chiefdoms support and are supported by culture of power (e.g. values of respect, honour, aggressiveness). This culture supports the interaction patterns that are required for maintaining the chiefdom structure in place (e.g. leaders dictate and subordinates obey). In turn, these culture and social structures support the occurrence of the same type of culture. In this study, Beck and Cowan (1996) relate each observable culture with very specific types of social reality (e.g. culture of power; culture of faith; culture of rationality).

These theories investigate the evolution of these triples of culture, society and environmental conditions. They explain that the occurrence of these triples appears to be ordered, that can be related to historical evolutions (e.g. first tribe then chiefdom then nation). When life or social conditions evolve, the social structure and culture evolve (often radically) from one level to the next. In Vanhée et al. (2013b), we propose a simulation model which replicates the relationship between social and environmental evolution. In this model, we assumed that culture was merged within social interactions but additional research for separating the two can be investigated.

This perspective raises important concerns for considering the design of culture for supporting coordination. First, it indicates that culture fits to limited environmental and social conditions. Second, it indicates that culture can evolve in order to fit better to the environment. In particular, it indicates the variety of possible social realities supported by culture and resulting social constructs.

Evolving culture provides a clear means for further increasing the autonomy of agent societies. It lets foreseeing the possibility for long-lasting systems, where agent societies, through cultural and social adaptiveness can adjust to a very wide range of environments.

8.4 Closing Remarks

As with many other solutions for coordinating agents (e.g. norms, organizations, protocols, nature-inspired approaches, cooperation), all started by an observation of living societies. From there, we discovered that individuals who share the same culture seem to manage better to coordinate with each other. Based on this observation, we tried to reproduce the same phenomenon within artificial societies, with the aim of being capable of obtaining similar improvements when coordinating artificial agents. And we did it!

In order to get to this result, we had to travel from limited informal descriptions provided by social sciences to computer science models. This was a long road. We first enriched social science theories, in order to get a better idea of the relationship between culture and coordination in human societies. Then, we had to find out what to model, determining what are the core aspects of culture involved in coordination and model these aspects. In order to be more confident that our model is coherent with human societies, we checked that this model credibly replicates the aspects of human societies we wanted. Then, we tried to better understand how cultures can be used for improving coordination, but here we faced an unexplored aspect of social sciences. We missed a crucial link for designing MASs: explaining how the known influence of culture on collective outcomes results from the known influence of culture on individual decisions. We overcame this issue through simulations and generative social science. We formulated a theory for explaining this relationship. This theory, in combination with our model, provided the basis for the final goal of this thesis: better understanding of how culture can be practically used for supporting flexible coordination.

In looking further away from the destination we reached, we can envision more clearly that culture can definitely help overcoming MAS challenges. Introducing artificial culture opens the gate towards another level of autonomy for multi-agent societies. Agents in these societies would be able of maintaining coherent interactions with each other by themselves, in spite of the wide variety of environments they may encounter. These societies would be able of interacting much more naturally with humans... And vice versa!

As teased in the introduction, we can now reconsider the issues I had there trying to coordinate in “the real world”. As we said there, when coordination failures occurred, we could hear the following sentence:

“Enfin, Loïs, un peu de bon sens !”
“Come on, Loïs, have a bit of common sense”

Therefore, was it a matter of common sense? As a cultural pun, yes and no! The English expression “common sense” encompasses two French expressions: “bon sens” and “sens commun”. However, in these two expressions, “common” differs in terms of meaning. The former means “common because external to us” (objective/rational²) while the latter means “common because everyone agrees upon” (consensual).

Attempting to solve the many coordination failures that arise from the “real world” as presented in the introduction can hardly be achieved with more “bon sens”. More “bon sens” would have lead us to endless arguments for finding the “optimal” answer (by the way, what can optimal mean?). While being useful for many decision problems, “bon

²More details about the meaning of “bon sens” in French is provided in Descartes’ “Discours de la Méthode”.

sens" is hard to apply alone for handling the real world: decision complexity increases too fast with regard to problem complexity. This is where values can play a role.

As a personal experience, it was much more effective to have a bit more of "sens commun". This "sens commun" provides shared general abstract *cultural values*, we agree on *what* to agree on (i.e. what is important for us, what do we keep in mind when considering things) and *how* to agree on what we want to agree on (i.e. how to interact with each other, what is important to keep in mind when we interact with each other). These two points help a lot for determining what we want and how we prefer to interact and coordinate in order to determine what we want.

Back on the example of my girlfriend, just by sharing the same importance we gave to the values of "cleanliness" versus "efficiency" and "following rules" versus "adapting to the current situation", we resolved so many conflicts at once! Based on these few values, we know what is important for both of us, we know how to interact with and what to expect from each other. Based on these indications, we can easily determine whether something is worth being pursued together. Then, when working together, values provide a general direction, helping to successfully interacting with each other when coordinating (e.g. building a detailed plan with sharp deadlines? Dividing abstract goals? Shall we do it together even if it is less efficient?). Plus, these few values offered a strong basis for building a much more complete social world (e.g. common desires, norms, usual plans, mutual expectations) for handling more "recurrent" interactions.

A bit like Voltaire's "Candide", one of my favorite books from high-school, we explored the world, learning from different aspects of culture, different ways of life. This exploration finally revealed a piece of fertile land, that can be cultivated for building concrete applications: culture can be used as a tool for supporting flexible coordination. Now, it is time for making use of this land: let's cultivate culture. As a borrow from Voltaire, the last sentence of his book is a reply of Candide to his former master Pangloss:

"Cela est bien dit, répondit Candide, mais il faut cultiver notre jardin".



Bibliography

- Aldewereld, H. and Dignum, V. (2010). Operetta: Organization-Oriented Development Environment. In *Languages, Methodologies, and Development Tools for Multi-Agent Systems*, (pp. 1–18).
- Antunes, L. and Coelho, H. (1999). Decisions Based upon Multiple Values: The {BVG} Agent Architecture. In P. Barahona and J. J. O. Alferes (Eds.), *Progress in Artificial Intelligence: Proceedings 9th Portuguese Conference on Artificial Intelligence, Lecture Notes in Computer Science*, (pp. 297–311). Springer.
- Axelrod, R. (1997). The Dissemination of Culture: A Model with Local Convergence and Global Polarization. *Journal of Conflict Resolution*, 41(2), 203–226.
- Beck, D. E. and Cowan, C. (1996). *Spiral Dynamics: Mastering Values, Leadership and Change*. Wiley-Blackwell.
- Bernstein, D., Givan, R., Immerman, N., and Zilberstein, S. (2002). The complexity of decentralized control of Markov decision processes. *Mathematics of Operations Research*, 27(4), 819–840.
- Bicchieri, C. (2006). *The Grammar of Society – The Nature and Dynamics of Social Norms*. Cambridge University Press.
- Binmore, K. (2008). *Playing for Real: A Text on Game Theory*. Oxford University Press.
- Boella, G., Torre, L., and Verhagen, H. (2006). Introduction to normative multiagent systems.
- Borit, M., Vanhée, L., and Olsen, P. (2013). Theoretical Considerations for Enhancing Social Believability through Integrating Culturally Specific Trust-Building Cognitive Features in Non-Playing Characters. In *Social Believability in Games Workshop, the Conference for Advances in Computer Entertainment*. Twente.
- Borit, M., Vanhée, L., and Olsen, P. (2014a). Towards enhancing trustworthiness of socially interactive and culture aware robots. In *Culture Aware Robotics (CAR2014)*.

- Borit, M., Vanh e, L., and Olsen, P. (2014b). Understanding the Impact of Culture on Cognitive Trust-Building Processes: How to Increase the Social Influence of Virtual Autonomous Agents. In *International Workshop on Trust in Agent Societies TRUST2014*.
- Brafman, R. I. and Domshlak, C. (2013). Preference Handling    An Introductory Tutorial. *AI Magazine*, 30, 58–86.
- Bratman, M. (1992). Shared cooperative activity. *The Philosophical Review*, 101(2), 327–341.
- Castelfranchi, C. (2000). Formalizing the Informal? In R. Demolombe and R. Hilpinen (Eds.), *Proceedings of the 5th International Workshop on Deontic Logic in Computer Science*. ONERA, Toulouse: Informal Proceedings.
- Castelfranchi, C., Dignum, F., Jonker, C. M., and Treur, J. (2000). Deliberative Normative Agents: Principles and Architecture. *Intelligent Agents VI, LNAI 1757*, 364–378.
- Conte, R., Castelfranchi, C., and Miceli, M. (1991). Limits and Levels of Cooperation: Disentangling Various Types of Prosocial Interaction. In Y. Demazeau and J.-P. Muller (Eds.), *Decentralized A.I. 2: Proc. of the 2nd European Workshop on Modelling Autonomous Agents in a Multi-Agent World*, (pp. 147–157).
- Dastani, M. (2008). 2APL: a practical agent programming language.
- Dawkins, R. (1976). *The Selfish Gene*.
- Dechesne, F., Tosto, G. D., Dignum, V., and Dignum, F. (2012). No smoking here: values, norms and culture in multi-agent systems. In *Artificial Intelligence and Law*, (pp. 79–107).
- Degens, N., Hofstede, G.-J., Mascarenhas, Samuel Silva, A., Paiva, A., Kistler, F., Andr e, E., Swiderska, A., Krumhuber, E., Kappas, A., Hume, C., Hall, L., and Aylett, R. (2013). Traveller - Intercultural training with intelligent agents for young adults. In *Proc. 1st International Workshop on Intelligent Digital Games for Empowerment and Inclusion (IDGEI 2013) held in conjunction with the 8th Foundations of Digital Games (FDG 2013)*. Chania, Crete, Greece: ACM, SASDG Digital Library.
- Dignum, F. (1999). Autonomous agents with norms. *Artificial Intelligence and Law*, 7(1), 69–79.
- Dignum, F., Kinny, D., and Sonenberg, L. (2001). Motivational Attitudes of Agents: On Desires, Obligations, and Norms. In *CEEMAS*, (pp. 83–92).
- Dignum, V. and Dignum, F. (2010). Designing agent systems: state of the practice. *IJAOSE*, 4(3), 224–243.
- D’Iribarne, P. (1989). *La logique de l’honneur. Gestion des entreprises et traditions nationales*. Seuil.
- Dunin-Keplicz, B. and Verbrugge, R. (2010). *Teamwork in Multi-Agent Systems: A Formal Approach*, Vol. 54. John Wiley and Sons.

- Durfee, E. H. (2006). Distributed Problem Solving and Planning. In M. Luck, V. MaĀžík, O. Štěpánková, and R. Trappl (Eds.), *Multi-Agent Systems and Applications*, chap. 3, (pp. 118–149). Lecture Notes in Computer Science, Springer.
- Endrass, B., André, E., Rehm, M., and Nakano, Y. (2013). Investigating culture-related aspects of behavior for virtual characters. *Autonomous Agents and MultiAgent Systems*, 27(2), 277–304.
- Epstein, J. M. (1999). Agent-based computational models and generative social science. *Complexity*, 4(5), 41–60.
- Epstein, J. M. and Axtell, R. (1996). *Growing artificial societies*. The MIT Press.
- Friedman, B., Kahn, P.H., J., and Borning, A. (2006). Value Sensitive Design and information systems. In P. Zhang and D. Galetta (Eds.), *Human-computer interaction and management information systems: Foundations*, (pp. 348–372). Advances in management information systems series, M.E. Sharpe.
- Gibson, C. B. and McDaniel, D. M. (2010). Moving Beyond Conventional Wisdom: Advancements in Cross-Cultural Theories of Leadership, Conflict, and Teams. *Perspectives on Psychological Science*, 5(4), 450–462.
- Grosz, B. J. (1996). Collaborative Systems. *AI Magazine*, (pp. 67–85).
- Gruber, B. T. (2005). What is an Ontology? *International Journal*, 2005(27 May 2005), 1–11.
- Hampden-Turner, C. and Trompenaars, A. (1993). *The seven cultures of capitalism: value systems for creating wealth in the United States, Japan, Germany, France, Britain, Sweden, and the Netherlands*. Currency Doubleday.
- Hofstede, G., Dignum, F., Prada, R., Student, J., and Vanhée, L. (2014). Gender difference: the role of nature, nurture, social identity and self-organization. In *The 15th International Workshop on Multi-Agent-Based Simulation MABS2014*. Paris, France.
- Hofstede, G., Hofstede, G. J., and Minkov, M. (2010a). *Cultures and Organizations: Software of the Mind, Third Edition*. McGraw-Hill Professional.
- Hofstede, G. J., Jonker, C. M., and Verwaart, T. (2010b). Cultural Differentiation of Negotiating Agents. *Group Decision and Negotiation*, 21(1), 79–98.
- Horling, B. and Lesser, V. (2005). A survey of multi-agent organizational paradigms. *The Knowledge Engineering Review*, 19(04), 281.
- Hubner, J. F., Sichman, J. S., and Boissier, O. (2004). Using the MOISE+ for a Cooperative Framework of MAS Reorganisation. In A. L. C. Bazzan and S. Labidi (Eds.), *SBLA*, Vol. 3171 of *LNAI*, (pp. 506–515). Springer.
- Kemper, T. D. (2011). *Status, power and ritual interaction: a relational reading of Durkheim, Goffman and Collins*.
- Kroeber, A. L. and Kluckhohn, C. (1952). *Culture: A critical review of concepts and definitions*, Vol. 47. The Museum.

- Levesque, H. J., Nunes, H. T., and Cohen, R. (1990). On Acting Together. In *Proceedings of the eighth National conference on Artificial intelligence (AAAI'90)*, (pp. 94–99). AAAI Press.
- Malone, T. W. and Crowston, K. (1990). What is Coordination Theory and How Can It Help Design Cooperative Work Systems? In *CSCW 90*, (pp. 357–370). No. October.
- Marzougui, B. and Barkaoui, K. (2013). Interaction Protocols in Multi-Agent Systems based on Agent Petri Nets Model. *International Journal of Advanced . . .*, 4(7), 166–173.
- Maskin, E. (2002). Implementation theory. In R. Aumann and S. Hart (Eds.), *Science*, Vol. 1, (pp. 237–288). Elsevier.
- Maslow, A. H. (1943). A theory of human motivation.
- Mc Breen, J., Di Tosto, G., Dignum, F., and Hofstede, G. J. (2011). Linking Norms and Culture. *2011 Second International Conference on Culture and Computing*, (pp. 9–14).
- Miceli, M. and Castelfranchi, C. (1989). A Cognitive Approach to Values. *Journal for the Theory of Social Behaviour*, 19(2), 169–193.
- Mintzberg, H. (1979). *The structuring of organizations: a synthesis of the research*. Prentice-Hall.
- Mintzberg, H. (1980). Structure in 5's: A Synthesis of the Research on Organization Design. *Management Science*, 26(3), 322–341.
- Morgenstern, O. (1951). Prolegomena to a theory of organization. Tech. rep.
- Ng, S. I., Lee, J. A., and Soutar, G. N. (2007). Are Hofstede's and Schwartz's value frameworks congruent?
- Norman, T. and Long, D. (1996). Alarms: An implementation of motivated agency. *Intelligent Agents II Agent Theories, Architectures, and Languages*, (pp. 219–234).
- Oliehoek, F. A. (2012). Decentralized POMDPs. In M. Wiering, Marco and Otterlo (Ed.), *Reinforcement Learning: State of the Art*, (pp. 471–503). Springer Berlin Heidelberg.
- Omicini, A., Ricci, A., Viroli, M., Castelfranchi, C., and Tummolini, L. (2004). Coordination artifacts: environment-based coordination for intelligent agents. *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems, 2004. AAMAS 2004*.
- Panait, L. and Luke, S. (2005). Cooperative Multi-Agent Learning: The State of the Art.
- Panait, L. A. and Luke, S. (2004). Learning Ant Foraging Behaviors. In J. Pollack, M. Bedau, P. Husbands, T. Ikegami, and R. A. Watson (Eds.), *Artificial Life {XI} Ninth International Conference on the Simulation and Synthesis of Living Systems*, (pp. 575–580). The MIT Press.
- Rao, A. S. and Georgeff, M. P. (1995). BDI Agents: From Theory to Practice. In *ICMAS*, (pp. 312–319).

- Reynolds, C. W. (1987). Flocks, herds and schools: A distributed behavioral model.
- Rokeach, M. (1973). Rokeach Values Survey. In *The Nature of Human Values*.
- Schwartz, S. (1999). A Theory of Cultural Values and Some Implications for Work. *Applied Psychology*, 48(1), 23.
- Schwartz, S. (2006a). An Overview Basic Human Values: Theory, Methods, and Applications Introduction to the Values Theory. *Jerusalem Hebrew University*.
- Schwartz, S. H. (1992). Universals in the Content and Structure of Values: Theoretical Advances and Empirical Tests in 20 Countries. *Advances in Experimental Social Psychology*, 25(C), 1–65.
- Schwartz, S. H. (2006b). Basic Human Values : An Overview Basic Human Values : Theory , Methods , and Applications Introduction to the Values Theory. *Jerusalem Hebrew University*, 48, 49–71.
- Sichman, J., Boissier, O., and Hannuon, M. (2002). A model for the structural, functional, and deontic specification of organizations in multiagent systems. *Advances in Artificial Intelligence*, (pp. 118–128).
- Steffan, R. R. (1997). Mintzberg on Management.
- Stone, P. and Veloso, M. (2000). Multiagent systems: A survey from a machine learning perspective. *Autonomous Robots*, 8(3), 1–57.
- Tuomela, R. (2000). *Cooperation: A Philosophical Study*, philosophi ed. Springer.
- Vanhée, L., Aldewereld, H., and Dignum, F. (2011). Implementing Norms? In *2011 IEEE/WICACM International Conferences on Web Intelligence and Intelligent Agent Technology*, Vol. 3, (pp. 13–16). IEEE.
- Vanhée, L., Dignum, F., and Ferber, J. (2013a). Towards Simulating the Impact of National Culture on Organizations. In *MABS2013: 14th International Workshop on Multi-Agent-Based Simulation*, (p. 12). Saint Paul.
- Vanhée, L., Dignum, F., and Ferber, J. (2014a). Modeling Culturally-Influenced Decisions. In *15th International Workshop on Multi-Agent-Based Simulation MABS2014 at the 13th International Conference on Autonomous Agents and Multiagent Systems*. Paris, France.
- Vanhée, L., Dignum, F., and Ferber, J. (2014b). Robust Collaboration: Enriching Decisions with Abstract Preferences. In *Engineering Multi-Agent Systems EMAS2014*.
- Vanhée, L., Ferber, J., and Dignum, F. (2013b). Agent-Based Evolving Societies. In *The 9th Conference of the European Social Simulation Association*. Warsaw.
- Vanhée, L., Ferber, J., and Dignum, F. (2013c). Agent-Based Evolving Societies (short paper). In *The 12th International Conference on Autonomous Agents and Multiagent Systems*. Saint Paul, Minnesota.
- Von Wright, G. H. (1951). Deontic Logic. *Mind*, 60(237), 1–15.

- Waldstrom, C. (2001). Informal Networks in Organizations - A literature review.
- van der Weide, T. (2011). *Arguing to motivate decisions*. Ph.D. thesis, Utrecht Universiteit.
- Wooldridge, M. (2002). *Introduction to Multiagent Systems*. John Wiley & Sons, Inc.
- Wooldridge, M. and Jennings, N. R. (1999). The cooperative problem-solving process. *Journal of Logic and Computation*, 9(4), 563–592.
- Yeoh, W. and Yokoo, M. (2012). Distributed Problem Solving. *AI Magazine* 33(3), (pp. 53–65).



Summary

This thesis proposes a method for supporting flexible coordination in multi-agent systems (MASs). In other words, we aim at influencing societies of artificial agents such that they can handle complex or evolving environments and collective goals (e.g. robots providing an emergency support capable of handling various hazards, climatic conditions, status of victims).

Towards achieving this goal, we first investigated why in human societies, for which MASs can be seen as an “artificial” counterpart, humans manage to coordinate relatively flexibly comparatively with artificial agents in MASs. We discovered that culture is a key factor of this relative success. Briefly, when humans share a cultural background, they manage to coordinate more flexibly because they share a common idea about what “working together” means. Conversely, artificial agents miss this aspect, leading in turn to coordination failures that can be similar to cultural clashes.

This missing aspect raises our goal: we want to better understand how culture can be integrated within and used for coordinating artificial societies. This goal raises the following research question: (how) can human-like culture be used as a tool for supporting coordination in artificial societies? As a preliminary step for answering this question, we need first to answer this question: (how) can the influence of human-like culture be integrated within artificial societies? In turn, this question raises a third one to be answered first: how does culture influence coordination in human societies?

As a first step, we expand general theories of culture for conceptualizing its influence in the context of coordination. From a generic perspective, we explain that culture influences individual decisions that support matching expectations and coherent interaction patterns, leading in turn to (generally) better collective performance. From a more specific perspective, we specify how the core acknowledged patterns of the influence of culture (e.g. cultural importance given to power status, to rules) apply in the context of coordination (e.g. culture influences the likeliness that leaders are (made) responsible for making

decisions for subordinates vs. proposing alternatives).

As a second step, we study how to replicate human-like influences of culture on coordination within artificial societies. First, since culture is grounded within individual decisions, we investigate the core culturally-sensitive decision aspects that impact the most (flexible) coordination in human societies. We discover that values, what people consider as “good” or “important” (e.g. honesty, obedience, autonomy), constitute such an aspect by deeply supporting a wide range of (interaction-related) decisions. Then, for illustrating how to replicate influence of culture within artificial societies, we build an value-sensitive agent decision architecture that can make coordination-related decisions. Finally, we illustrate that our architecture can replicate the influence of culture on coordination through two simulations that replicate known coordination-related cultural phenomena.

As a third step, we study how human-like values can be used for supporting coordination in artificial societies. First, we investigate the range of coordination problems for which values can offer an operational means for supporting coordination. As in human societies, values are particularly adequate for problems with complex and dynamic environments, requiring agents to make coordination-related decisions. Then, towards concretely implementing values, we study the technical details to consider when using values for supporting flexible coordination (e.g. how to concretely design values and integrating them within decision processes).

As a summary, this thesis highlights that key aspects of the influence of culture on coordination (particularly, through the influence of culture on values) can be replicated within artificial societies. Furthermore, we show that this influence can be handled such that culture can be used as means for supporting flexible coordination in artificial societies.



Résumé

Cette thèse propose une méthode pour coordonner flexiblement des Systèmes Multi-Agents (SMA). Plus en détails, nous étudions comment influencer des agents artificiels afin que, collectivement, ils atteignent des objectifs complexes et/ou dynamiques dans des environnements eux-aussi complexes et dynamiques (ex : un groupe de robots pour secourir les victimes lors d'un désastre, qui peut s'adapter à une grande variété de dangers, conditions climatiques, état des victimes).

Dans ce but, nous avons d'abord étudié pourquoi, dans les sociétés humaines, les humains parviennent à coordonner relativement flexiblement mais pas leurs contreparties artificielles (agents des SMA). Cette opposition peut être grandement expliquée à l'aide d'un facteur clef : la culture. Les humains qui partagent un même bagage culturel se coordonnent flexiblement plus facilement, car ils ont une idée commune de ce que "travailler ensemble" veut dire. A contrario, les agents n'ont pas ce bagage et leurs échecs pour travailler ensemble s'apparentent souvent à des chocs culturels.

Ainsi, notre objectif consiste à répondre à la question suivante : peut-on utiliser une culture semblable à celle des humains comme un outil pour coordonner les SMA (et si oui, comment) ? Pour répondre à cette question, il nous faut d'abord expliquer : comment intégrer une culture semblable à celle des humains dans un SMA ? Cette seconde question en soulève une troisième à étudier en premier : comment est-ce que la culture influence la manière dont la coordination se passe dans les sociétés humaines ?

1- Nous montrons que de manière générale, la culture influence les décisions individuelles prises en situation d'interaction (ex : au travers d'attentes, de manière d'agir et de raisonner). Cette influence mène à l'occurrence de schémas d'interaction abstraits, récurrents et cohérents, qui, généralement, améliorent la performance collective. Ensuite, nous spécifions comment les principaux mécanismes de l'influence connue de la culture (ex : importance culturelle accordée au pouvoir, aux règles) s'appliquent spécifiquement en situation de coordination (ex : la culture influence si les dirigeants donnent des ordres vs. des propositions à leurs subordonnés).

2-Nous montrons comment répliquer les mécanismes l'influence de la culture sur la coordination dans les SMA. Tout d'abord, puisque la culture est fondée dans les décisions individuelles, nous mettons en avant un mécanisme de décision humain clef qui, à la fois, est sensible à la culture et influence la coordination. Ce mécanisme se trouve dans les valeurs, ce que les gens considèrent comme "bien" ou "important" (ex : honnêteté, discipline, autonomie). Ensuite, nous intégrons ces valeurs dans une architecture agent capable de prendre des décisions en situation de coordination. Enfin, nous illustrons que notre architecture peut en effet reproduire l'influence de la culture sur la coordination à travers de deux simulations qui répliquent des phénomènes culturels en situation de coordination connus.

3-Nous étudions comment ces valeurs, inspirées des valeurs humaines, peuvent être utilisées coordonner des SMA. Tout d'abord, nous étudions pour quels problèmes les valeurs offrent un moyen opérationnel pour soutenir la coordination. A l'instar des sociétés humaines, les valeurs sont particulièrement offrent un haut niveau de flexibilité, quand les agents doivent raisonner eux-même pour établir une coordination. Puis, nous étudions les détails techniques à considérer pour utiliser en pratique des valeurs pour coordonner flexiblement des SMA (ex : quelles valeurs choisir ? Comment les représenter ?).

En résumé, cette thèse met en évidence que les principaux mécanismes de l'influence de la culture sur la coordination (en particulier, grâce à l'influence de la culture sur les valeurs) peuvent être répliquées au sein des SMA. De plus, nous montrons que ces mécanismes peuvent être manipulés dans le but de coordonner des SMA.



Samenvatting

Verschillende applicaties vereisen samenwerking van computationele entiteiten. Denk bijvoorbeeld aan teams van drones die slachtoffers zoeken in rampgebieden, of zogeheten smart homes die informatie verzamelen ter ondersteuning van patiënten. Een oplossing voor het omgaan met dergelijke applicaties is om de programma's te zien als "individuen" of "personen"; we noemen ze dan "agenten". Met dit perspectief kunnen we agenten samen laten werken zoals we mensen laten samenwerken. Bijvoorbeeld door het maken van (artificiële) organisaties of regels.

Het is nog steeds een uitdaging om agenten flexibel te laten samenwerken wanneer het systeem verscheidene taken heeft en moet omgaan met veel verschillende situaties. In een menselijke maatschappij lijkt dit veel minder een probleem te zijn. Daarom hebben we gekeken naar menselijke maatschappijen, en ons afgevraagd "Waarom kunnen mensen goed samenwerken en agenten niet? Waar ontbreekt het aan bij agenten?"

Agenten hebben geen goed gemeenschappelijk beeld over wat samenwerken betekent. Daarom kunnen ze zich niet goed aanpassen aan nieuwe situaties, wat essentieel is voor flexibiliteit. We observeren in menselijke maatschappijen dat cultuur een zeer belangrijke rol speelt bij het bepalen wat mensen onder samenwerken verstaan. Daarom hebben we de notie van cultuur aan (artificiële) agenten gegeven, in de hoop dat ze dan beter samen kunnen werken.

Ten eerste hebben we aspecten van menselijke cultuur gerepliceerd in artificiële agenten. Ten tweede hebben we simulaties gemaakt die basis culturele fenomenen reproduceren die we kunnen observeren in menselijke maatschappijen. Ten derde hebben we onderzocht hoe cultuur gebruikt kan worden om de samenwerking tussen agenten te verbeteren.



Acknowledgements, Dankwoord, Remerciements

This thesis is a part of a longer life travel that involved numerous people who share a part of responsibility in the completion this text. As a token of gratitude for their support, I want to thank them all here.

First of all, I want to thank my work partners, with whom I directly worked during this thesis, starting with my supervisors: Frank Dignum, Jacques Ferber and John-Jules Charles Meyer. No need for details here, they already know how much I owe them for their involvement throughout my thesis. I also want to thank the reading committee: Olivier Boissier, Christiano Castelfranchi, Amal El Fallah Seghrouchni, Emma Norling and Jan Treur. They dedicated a substantial amount of time to evaluate my thesis and indicating how to improve it. I also want to thank the partners with whom I shared several work projects: Éric Bourreau, Gert-Jan Hofstede, Jillian Student, Melania Borit and Rui Prada.

I want to thank my funders. ENS Cachan funded an initial year of “free work” in Utrecht and provided the means for receiving a PhD grant from the University of Montpellier 2. I also want to thank our department and in particular its head, Linda van der Gaag. In spite of the tight budget for running the department, she accepted to fund the last months of my thesis writing period. This offer allowed overcoming a difficult administrative issues that would have prevented the thesis from being completed. Finally, I want to thank Frank Dignum who connected me with projects that I could perform along my thesis, limiting the stress caused by financial difficulties raised by working in a more expensive country.

I want to thank people who directly supported me during this PhD. I want to thank Nicolas Serrurier, one of the secretary from the LIRMM. In many occasions, he provided the necessary help for handling the numerous formal obligations raised by the French administration. Details about his support are better skipped here, but there is little doubt that my PhD could have failed without him. I also want to thank Melania Borit who provided extensive help throughout this PhD, by reading multiple times through my articles and my thesis, sharply detecting how to improve the quality of the text and raising pertinent (and sometimes annoyingly tough) questions. Furthermore, she gave me a lot of advice for improving my English, helping me to progress over years, keeping me looking back to my former texts, which made me think “gosh, did I write that?”. Without any doubt, both of them easily saved me weeks (if not months) of worktime that I could invest instead in research.

I want to thank people who indirectly lead to the success of this PhD. I want first to thank my mentors, or inspirational figures who served as role models I used to build on my own character and identity. At the core, I want to bow to Alain Pilfer, my first Kendo senseï. He taught me the values of hard work, dedication, commitment, search for excellence, self-improvement and going beyond my limits. I also want to thank Jay Francis, my squash teacher, who completed this learning process by teaching me the potential benefits of using quick and dirty rules that fit the job. More professionally-related, I want to thank Matthieu Exbrayat and Nathalie Bertrand, two former supervisors with whom I discovered the enjoyment of working for helping people. Likewise, I want to thank Jean-Jacques Lacrampe and Sjoerd Timmer, who taught me about the pleasure of practicing science, just for the sake of doing it. Finally, I want to thank Henri Thuillier and Ioan Todinca, who believed in my capabilities and showed me the way to the ENS.

I also want to thank all the others who were of enjoyable company during this trip: family, friends, colleagues. Here is a (non-exhaustive) list: Bas, Benjamin, Bernadette, Bernard, Catherine, Charles, Cheah, David, Eric, Fabien, Florian, Françoise Gennaro, Jan, Janneke, Jean-Baptiste, Joost, "de andere" Joost, Karine, Luora, Krzysztof, Margaux, Maaïke, Marieke, Max, Michal, Mohamed, Roger, Roland, Sébastien, Shos, Sjoerd, 'Tjan, Tom... without forgetting kendo and squash partners, angels, demons, pirate crews and field mechanics. Amongst these people, I met some of the best friends I have ever had.

Last but far from the least, I want to thank Solenne. She was always there, caring about me and supporting me, in spite of the distance and my lack of time that separated us during this thesis. I am glad that we are now finally reunited.

To all of them, merci !

Paris,
22/04/2015

Loïs Vanhée

SIKS Dissertation Series

1998 |

- 1998-1 | **Johan van den Akker (CWI)** DEGAS - An Active, Temporal Database of Autonomous Objects
1998-2 | **Floris Wiesman (UM)** Information Retrieval by Graphically Browsing Meta-Information
1998-3 | **Ans Steuten (TUD)** A Contribution to the Linguistic Analysis of Business Conversations within the Language/Action Perspective
1998-4 | **Dennis Breuker (UM)** Memory versus Search in Games
1998-5 | **E.W.Oskamp (RUL)** Computerondersteuning bij Straftoemeting

1999 |

- 1999-1 | **Mark Sloof (VU)** Physiology of Quality Change Modelling: Automated modelling of Quality Change of Agricultural Products
1999-2 | **Rob Potharst (EUR)** Classification using decision trees and neural nets
1999-3 | **Don Beal (UM)** The Nature of Minimax Search
1999-4 | **Jacques Penders (UM)** The practical Art of Moving Physical Objects
1999-5 | **Aldo de Moor (KUB)** Empowering Communities: A Method for the Legitimate User-Driven Specification of Network Information Systems
1999-6 | **Niek J.E. Wijngaards (VU)** Re-design of compositional systems
1999-7 | **David Spelt (UT)** Verification support for object database design
1999-8 | **Jacques H.J. Lenting (UM)** Informed Gambling: Conception and Analysis of a Multi-Agent Mechanism for Discrete Reallocation.

2000 |

- 2000-1 | **Frank Niessink (VU)** Perspectives on Improving Software Maintenance
2000-2 | **Koen Holtman (TUE)** Prototyping of CMS Storage Management
2000-3 | **Carolien M.T. Metselaar (UVA)** Sociaal-organisatorische gevolgen van kennistechnologie; een procesbenadering en actorperspectief.
2000-4 | **Geert de Haan (VU)** ETAG, A Formal Model of Competence Knowledge for User Interface Design
2000-5 | **Ruud van der Pol (UM)** Knowledge-based Query Formulation in Information Retrieval.
2000-6 | **Rogier van Eijk (UU)** Programming Languages for Agent Communication
2000-7 | **Niels Peek (UU)** Decision-theoretic Planning of Clinical Patient Management
2000-8 | **Veerie Coupé; (EUR)** Sensitivity Analysis of Decision-Theoretic Networks
2000-9 | **Florian Waas (CWI)** Principles of Probabilistic Query Optimization
2000-10 | **Niels Nes (CWI)** Image Database Management System Design Considerations, Algorithms and Architecture
2000-11 | **Jonas Karlsson (CWI)** Scalable Distributed Data Structures for Database Management

2001 |

- 2001-1 | **Silja Renooij (UU)** Qualitative Approaches to Quantifying Probabilistic Networks
2001-2 | **Koen Hindriks (UU)** Agent Programming Languages: Programming with Mental Models
2001-3 | **Maarten van Someren (UvA)** Learning as problem solving
2001-4 | **Evgueni Smirnov (UM)** Conjunctive and Disjunctive Version Spaces with Instance-Based Boundary Sets

- 2001-5 | **Jacco van Ossenbruggen (VU)** Processing Structured Hypermedia: A Matter of Style
2001-6 | **Martijn van Welie (VU)** Task-based User Interface Design
2001-7 | **Bastiaan Schonhage (VU)** Diva: Architectural Perspectives on Information Visualization
2001-8 | **Pascal van Eck (VU)** A Compositional Semantic Structure for Multi-Agent Systems Dynamics.
2001-9 | **Pieter Jan 't Hoen (RUL)** Towards Distributed Development of Large Object-Oriented Models, Views of Packages as Classes
2001-10 | **Maarten Sierhuis (UvA)** Modeling and Simulating Work Practice BRAHMS: a multiagent modeling and simulation language for work practice analysis and design
2001-11 | **Tom M. van Engers (VUA)** Knowledge Management: The Role of Mental Models in Business Systems Design

2002 |

- 2002-01 | **Nico Lassing (VU)** Architecture-Level Modifiability Analysis
2002-02 | **Roelof van Zwol (UT)** Modelling and searching web-based document collections
2002-03 | **Henk Ernst Blok (UT)** Database Optimization Aspects for Information Retrieval
2002-04 | **Juan Roberto Castelo Valdueza (UU)** The Discrete Acyclic Digraph Markov Model in Data Mining
2002-05 | **Radu Serban (VU)** The Private Cyberspace Modeling Electronic Environments inhabited by Privacy-concerned Agents
2002-06 | **Laurens Mommers (UL)** Applied legal epistemology; Building a knowledge-based ontology of the legal domain
2002-07 | **Peter Boncz (CWI)** Monet: A Next-Generation DBMS Kernel For Query-Intensive Applications
2002-08 | **Jaap Gordijn (VU)** Value Based Requirements Engineering: Exploring Innovative E-Commerce Ideas
2002-09 | **Willem-Jan van den Heuvel(KUB)** Integrating Modern Business Applications with Objectified Legacy Systems
2002-10 | **Brian Sheppard (UM)** Towards Perfect Play of Scrabble
2002-11 | **Wouter C.A. Wijngaards (VU)** Agent Based Modelling of Dynamics: Biological and Organisational Applications
2002-12 | **Albrecht Schmidt (Uva)** Processing XML in Database Systems
2002-13 | **Hongjing Wu (TUE)** A Reference Architecture for Adaptive Hypermedia Applications
2002-14 | **Wieke de Vries (UU)** Agent Interaction: Abstract Approaches to Modelling, Programming and Verifying Multi-Agent Systems
2002-15 | **Rik Eshuis (UT)** Semantics and Verification of UML Activity Diagrams for Workflow Modelling
2002-16 | **Pieter van Langen (VU)** The Anatomy of Design: Foundations, Models and Applications
2002-17 | **Stefan Manegold (UVA)** Understanding, Modeling, and Improving Main-Memory Database Performance

2003 |

- 2003-01 | **Heiner Stuckenschmidt (VU)** Ontology-Based Information Sharing in Weakly Structured Environments
2003-02 | **Jan Broersen (VU)** Modal Action Logics for Reasoning About Reactive Systems
2003-03 | **Martijn Schuemie (TUD)** Human-Computer Interaction and Presence in Virtual Reality Exposure Therapy
2003-04 | **Milan Petkovic (UT)** Content-Based Video Retrieval Supported by Database Technology

- 2003-05 | **Jos Lehmann (UVA)** Causation in Artificial Intelligence and Law - A modelling approach
- 2003-06 | **Boris van Schooten (UT)** Development and specification of virtual environments
- 2003-07 | **Machiel Jansen (UvA)** Formal Explorations of Knowledge Intensive Tasks
- 2003-08 | **Yongping Ran (UM)** Repair Based Scheduling
- 2003-09 | **Rens Kortmann (UM)** The resolution of visually guided behaviour
- 2003-10 | **Andreas Lincke (UvT)** Electronic Business Negotiation: Some experimental studies on the interaction between medium, innovation context and culture
- 2003-11 | **Simon Keizer (UT)** Reasoning under Uncertainty in Natural Language Dialogue using Bayesian Networks
- 2003-12 | **Roeland Ordelman (UT)** Dutch speech recognition in multimedia information retrieval
- 2003-13 | **Jeroen Donkers (UM)** Nosce Hostem - Searching with Opponent Models
- 2003-14 | **Stijn Hoppenbrouwers (KUN)** Freezing Language: Conceptualisation Processes across ICT-Supported Organisations
- 2003-15 | **Mathijs de Weerd (TUD)** Plan Merging in Multi-Agent Systems
- 2003-16 | **Menzo Windhouwer (CWI)** Feature Grammar Systems - Incremental Maintenance of Indexes to Digital Media Warehouses
- 2003-17 | **David Jansen (UT)** Extensions of Statecharts with Probability, Time, and Stochastic Timing
- 2003-18 | **Levente Kocsis (UM)** Learning Search Decisions
- 2004 |**
- 2004-01 | **Virginia Dignum (UU)** A Model for Organizational Interaction: Based on Agents, Founded in Logic
- 2004-02 | **Lai Xu (UvT)** Monitoring Multi-party Contracts for E-business
- 2004-03 | **Perry Groot (VU)** A Theoretical and Empirical Analysis of Approximation in Symbolic Problem Solving
- 2004-04 | **Chris van Aart (UVA)** Organizational Principles for Multi-Agent Architectures
- 2004-05 | **Vlára Popova (EUR)** Knowledge discovery and monotonicity
- 2004-06 | **Bart-Jan Hommes (TUD)** The Evaluation of Business Process Modeling Techniques
- 2004-07 | **Elise Boltjes (UM)** Voorbeeldig onderwijs; voorbeeldgestuurd onderwijs, een opstap naar abstract denken, vooral voor meisjes
- 2004-08 | **Joop Verbeek (UM)** Politie en de Nieuwe Internationale Informatiemarkt, Grensregionale politie en digitale expertise
- 2004-09 | **Martin Caminada (VU)** For the Sake of the Argument; explorations into argument-based reasoning
- 2004-10 | **Suzanne Kabel (UVA)** Knowledge-rich indexing of learning-objects
- 2004-11 | **Michel Klein (VU)** Change Management for Distributed Ontologies
- 2004-12 | **The Duy Bui (UT)** Creating emotions and facial expressions for embodied agents
- 2004-13 | **Wojciech Jamroga (UT)** Using Multiple Models of Reality: On Agents who Know how to Play
- 2004-14 | **Paul Harrenstein (UU)** Logic in Conflict. Logical Explorations in Strategic Equilibrium
- 2004-15 | **Arno Knobbe (UU)** Multi-Relational Data Mining
- 2004-16 | **Federico Divina (VU)** Hybrid Genetic Relational Search for Inductive Learning
- 2004-17 | **Mark Winands (UM)** Informed Search in Complex Games
- 2004-18 | **Vania Bessa Machado (UvA)** Supporting the Construction of Qualitative Knowledge Models
- 2004-19 | **Thijs Westerveld (UT)** Using generative probabilistic models for multimedia retrieval
- 2004-20 | **Madelon Evers (Nyenrode)** Learning from Design: facilitating multidisciplinary design teams
- 2005 |**
- 2005-01 | **Floor Verdenius (UVA)** Methodological Aspects of Designing Induction-Based Applications
- 2005-02 | **Erik van der Werf (UM)** AI techniques for the game of Go
- 2005-03 | **Franc Grootjen (RUN)** A Pragmatic Approach to the Conceptualisation of Language
- 2005-04 | **Nirvana Meratnia (UT)** Towards Database Support for Moving Object data
- 2005-05 | **Gabriel Infante-Lopez (UVA)** Two-Level Probabilistic Grammars for Natural Language Parsing
- 2005-06 | **Pieter Spronck (UM)** Adaptive Game AI
- 2005-07 | **Flavius Frasinca (TUE)** Hypermedia Presentation Generation for Semantic Web Information Systems
- 2005-08 | **Richard Vdovjak (TUE)** A Model-driven Approach for Building Distributed Ontology-based Web Applications
- 2005-09 | **Jeen Broekstra (VU)** Storage, Querying and Inferencing for Semantic Web Languages
- 2005-10 | **Anders Bouwer (UVA)** Explaining Behaviour: Using Qualitative Simulation in Interactive Learning Environments
- 2005-11 | **Elth Ogston (VU)** Agent Based Matchmaking and Clustering - A Decentralized Approach to Search
- 2005-12 | **Csaba Boer (EUR)** Distributed Simulation in Industry
- 2005-13 | **Fred Hamburg (UL)** Een Computermodel voor het Ondersteunen van Euthanasiebeslissingen
- 2005-14 | **Borys Omelayenko (VU)** Web-Service configuration on the Semantic Web; Exploring how semantics meets pragmatics
- 2005-15 | **Tibor Bosse (VU)** Analysis of the Dynamics of Cognitive Processes
- 2005-16 | **Joris Graaumanns (UU)** Usability of XML Query Languages
- 2005-17 | **Boris Shishkov (TUD)** Software Specification Based on Re-usable Business Components
- 2005-18 | **Danielle Sent (UU)** Test-selection strategies for probabilistic networks
- 2005-19 | **Michel van Dartel (UM)** Situated Representation
- 2005-20 | **Cristina Coteanu (UL)** Cyber Consumer Law, State of the Art and Perspectives
- 2005-21 | **Wijnand Derks (UT)** Improving Concurrency and Recovery in Database Systems by Exploiting Application Semantics
- 2006 |**
- 2006-01 | **Samuil Angelov (TUE)** Foundations of B2B Electronic Contracting
- 2006-02 | **Cristina Chisalita (VU)** Contextual issues in the design and use of information technology in organizations
- 2006-03 | **Noor Christoph (UVA)** The role of metacognitive skills in learning to solve problems
- 2006-04 | **Marta Sabou (VU)** Building Web Service Ontologies
- 2006-05 | **Cees Pierik (UU)** Validation Techniques for Object-Oriented Proof Outlines
- 2006-06 | **Ziv Baida (VU)** Software-aided Service Bundling - Intelligent Methods & Tools for Graphical Service Modeling
- 2006-07 | **Marko Smiljanic (UT)** XML schema matching - balancing efficiency and effectiveness by means of clustering
- 2006-08 | **Eelco Herder (UT)** Forward, Back and Home Again - Analyzing User Behavior on the Web

- 2006-09 | **Mohamed Wahdan (UM)** Automatic Formulation of the Auditor's Opinion
- 2006-10 | **Ronny Siebes (VU)** Semantic Routing in Peer-to-Peer Systems
- 2006-11 | **Joeri van Ruth (UT)** Flattening Queries over Nested Data Types
- 2006-12 | **Bert Bongers (VU)** Interactivation - Towards an e-cology of people, our technological environment, and the arts
- 2006-13 | **Henk-Jan Lebbink (UU)** Dialogue and Decision Games for Information Exchanging Agents
- 2006-14 | **Johan Hoorn (VU)** Software Requirements: Update, Upgrade, Redesign - towards a Theory of Requirements Change
- 2006-15 | **Rainer Malik (UU)** CONAN: Text Mining in the Biomedical Domain
- 2006-16 | **Carsten Riggelsen (UU)** Approximation Methods for Efficient Learning of Bayesian Networks
- 2006-17 | **Stacey Nagata (UU)** User Assistance for Multi-tasking with Interruptions on a Mobile Device
- 2006-18 | **Valentin Zhizhkun (UVA)** Graph transformation for Natural Language Processing
- 2006-19 | **Birna van Riemsdijk (UU)** Cognitive Agent Programming: A Semantic Approach
- 2006-20 | **Marina Velikova (UvT)** Monotone models for prediction in data mining
- 2006-21 | **Bas van Gils (RUN)** Aptness on the Web
- 2006-22 | **Paul de Vrieze (RUN)** Fundamentals of Adaptive Personalisation
- 2006-23 | **Ion Juvina (UU)** Development of Cognitive Model for Navigating on the Web
- 2006-24 | **Laura Hollink (VU)** Semantic Annotation for Retrieval of Visual Resources
- 2006-25 | **Madalina Drugan (UU)** Conditional log-likelihood MDL and Evolutionary MCMC
- 2006-26 | **Vojkan Mihajlovic (UT)** Score Region Algebra: A Flexible Framework for Structured Information Retrieval
- 2006-27 | **Stefano Bocconi (CWI)** Vox Populi: generating video documentaries from semantically annotated media repositories
- 2006-28 | **Borkur Sigurbjornsson (UVA)** Focused Information Access using XML Element Retrieval
- 2007 |**
- 2007-01 | **Kees Leune (UvT)** Access Control and Service-Oriented Architectures
- 2007-02 | **Wouter Teepe (RUG)** Reconciling Information Exchange and Confidentiality: A Formal Approach
- 2007-03 | **Peter Mika (VU)** Social Networks and the Semantic Web
- 2007-04 | **Jurriaan van Diggelen (UU)** Achieving Semantic Interoperability in Multi-agent Systems: a dialogue-based approach
- 2007-05 | **Bart Schermer (UL)** Software Agents, Surveillance, and the Right to Privacy: a Legislative Framework for Agent-enabled Surveillance
- 2007-06 | **Gilad Mishne (UVA)** Applied Text Analytics for Blogs
- 2007-07 | **Natasa Jovanovic' (UT)** To Whom It May Concern - Addressee Identification in Face-to-Face Meetings
- 2007-08 | **Mark Hoogendoorn (VU)** Modeling of Change in Multi-Agent Organizations
- 2007-09 | **David Mobach (VU)** Agent-Based Mediated Service Negotiation
- 2007-10 | **Huib Aldewereld (UU)** Autonomy vs. Conformity: an Institutional Perspective on Norms and Protocols
- 2007-11 | **Natalia Stash (TUE)** Incorporating Cognitive/Learning Styles in a General-Purpose Adaptive Hypermedia System
- 2007-12 | **Marcel van Gerven (RUN)** Bayesian Networks for Clinical Decision Support: A Rational Approach to Dynamic Decision-Making under Uncertainty
- 2007-13 | **Rutger Rienks (UT)** Meetings in Smart Environments: Implications of Progressing Technology
- 2007-14 | **Niek Bergboer (UM)** Context-Based Image Analysis
- 2007-15 | **Joyca Lacroix (UM)** NIM: a Situated Computational Memory Model
- 2007-16 | **Davide Grossi (UU)** Designing Invisible Handcuffs. Formal investigations in Institutions and Organizations for Multi-agent Systems
- 2007-17 | **Theodore Charitos (UU)** Reasoning with Dynamic Networks in Practice
- 2007-18 | **Bart Orriens (UvT)** On the development and management of adaptive business collaborations
- 2007-19 | **David Levy (UM)** Intimate relationships with artificial partners
- 2007-20 | **Slinger Jansen (UU)** Customer Configuration Updating in a Software Supply Network
- 2007-21 | **Karianne Vermaas (UU)** Fast diffusion and broadening use: A research on residential adoption and usage of broadband internet in the Netherlands between 2001 and 2005
- 2007-22 | **Zlatko Zlatev (UT)** Goal-oriented design of value and process models from patterns
- 2007-23 | **Peter Barna (TUE)** Specification of Application Logic in Web Information Systems
- 2007-24 | **Georgina Ramı́rez Camps (CWI)** Structural Features in XML Retrieval
- 2007-25 | **Joost Schalken (VU)** Empirical Investigations in Software Process Improvement
- 2008 |**
- 2008-01 | **Katalin Boer-Sorbı́n (EUR)** Agent-Based Simulation of Financial Markets: A modular, continuous-time approach
- 2008-02 | **Alexei Sharpanskykh (VU)** On Computer-Aided Methods for Modeling and Analysis of Organizations
- 2008-03 | **Vera Hollink (UVA)** Optimizing hierarchical menus: a usage-based approach
- 2008-04 | **Ander de Keijzer (UT)** Management of Uncertain Data - towards unattended integration
- 2008-05 | **Bela Mutschler (UT)** Modeling and simulating causal dependencies on process-aware information systems from a cost perspective
- 2008-06 | **Arjen Hommersom (RUN)** On the Application of Formal Methods to Clinical Guidelines, an Artificial Intelligence Perspective
- 2008-07 | **Peter van Rosmalen (OU)** Supporting the tutor in the design and support of adaptive e-learning
- 2008-08 | **Janneke Bolt (UU)** Bayesian Networks: Aspects of Approximate Inference
- 2008-09 | **Christof van Nimwegen (UU)** The paradox of the guided user: assistance can be counter-effective
- 2008-10 | **Wauter Bosma (UT)** Discourse oriented summarization
- 2008-11 | **Vera Kartseva (VU)** Designing Controls for Network Organizations: A Value-Based Approach
- 2008-12 | **Jozsef Farkas (RUN)** A Semiotically Oriented Cognitive Model of Knowledge Representation
- 2008-13 | **Caterina Carraciolo (UVA)** Topic Driven Access to Scientific Handbooks
- 2008-14 | **Arthur van Bunnigen (UT)** Context-Aware Querying: Better Answers with Less Effort
- 2008-15 | **Martijn van Otterlo (UT)** The Logic of Adaptive Behavior: Knowledge Representation and Algorithms for the Markov Decision Process Framework in First-Order Domains.

- 2008-16 | **Henriette van Vugt (VU)** Embodied agents from a user's perspective
- 2008-17 | **Martin Op 't Land (TUD)** Applying Architecture and Ontology to the Splitting and Allying of Enterprises
- 2008-18 | **Guido de Croon (UM)** Adaptive Active Vision
- 2008-19 | **Henning Rode (UT)** From Document to Entity Retrieval: Improving Precision and Performance of Focused Text Search
- 2008-20 | **Rex Arendsen (UVA)** Geen bericht, goed bericht. Een onderzoek naar de effecten van de introductie van elektronisch berichtenverkeer met de overheid op de administratieve lasten van bedrijven.
- 2008-21 | **Krisztian Balog (UVA)** People Search in the Enterprise
- 2008-22 | **Henk Koning (UU)** Communication of IT-Architecture
- 2008-23 | **Stefan Visscher (UU)** Bayesian network models for the management of ventilator-associated pneumonia
- 2008-24 | **Zharko Aleksovski (VU)** Using background knowledge in ontology matching
- 2008-25 | **Geert Jonker (UU)** Efficient and Equitable Exchange in Air Traffic Management Plan Repair using Spender-signed Currency
- 2008-26 | **Marijn Huijbregts (UT)** Segmentation, Diarization and Speech Transcription: Surprise Data Unraveled
- 2008-27 | **Hubert Vogten (OU)** Design and Implementation Strategies for IMS Learning Design
- 2008-28 | **Ildiko Flesch (RUN)** On the Use of Independence Relations in Bayesian Networks
- 2008-29 | **Dennis Reidsma (UT)** Annotations and Subjective Machines - Of Annotators, Embodied Agents, Users, and Other Humans
- 2008-30 | **Wouter van Atteveldt (VU)** Semantic Network Analysis: Techniques for Extracting, Representing and Querying Media Content
- 2008-31 | **Loes Braun (UM)** Pro-Active Medical Information Retrieval
- 2008-32 | **Trung H. Bui (UT)** Toward Affective Dialogue Management using Partially Observable Markov Decision Processes
- 2008-33 | **Frank Terpstra (UVA)** Scientific Workflow Design: theoretical and practical issues
- 2008-34 | **Jeroen de Knijf (UU)** Studies in Frequent Tree Mining
- 2008-35 | **Ben Torben Nielsen (UvT)** Dendritic morphologies: function shapes structure
- 2009 |**
- 2009-01 | **Rasa Jurgelenaite (RUN)** Symmetric Causal Independence Models
- 2009-02 | **Willem Robert van Hage (VU)** Evaluating Ontology-Alignment Techniques
- 2009-03 | **Hans Stol (UvT)** A Framework for Evidence-based Policy Making Using IT
- 2009-04 | **Josephine Nabukenya (RUN)** Improving the Quality of Organisational Policy Making using Collaboration Engineering
- 2009-05 | **Sietse Overbeek (RUN)** Bridging Supply and Demand for Knowledge Intensive Tasks - Based on Knowledge, Cognition, and Quality
- 2009-06 | **Muhammad Subianto (UU)** Understanding Classification
- 2009-07 | **Ronald Poppe (UT)** Discriminative Vision-Based Recovery and Recognition of Human Motion
- 2009-08 | **Volker Nannen (VU)** Evolutionary Agent-Based Policy Analysis in Dynamic Environments
- 2009-09 | **Benjamin Kanagwa (RUN)** Design, Discovery and Construction of Service-oriented Systems
- 2009-10 | **Jan Wielemaker (UVA)** Logic programming for knowledge-intensive interactive applications
- 2009-11 | **Alexander Boer (UVA)** Legal Theory, Sources of Law & the Semantic Web
- 2009-12 | **Peter Massuthe (TUE, Humboldt-Universitaet zu Berlin)** Operating Guidelines for Services
- 2009-13 | **Steven de Jong (UM)** Fairness in Multi-Agent Systems
- 2009-14 | **Maksym Korotkiy (VU)** From ontology-enabled services to service-enabled ontologies (making ontologies work in e-science with ONTO-SOA)
- 2009-15 | **Rinke Hoekstra (UVA)** Ontology Representation - Design Patterns and Ontologies that Make Sense
- 2009-16 | **Fritz Reul (UvT)** New Architectures in Computer Chess
- 2009-17 | **Laurens van der Maaten (UvT)** Feature Extraction from Visual Data
- 2009-18 | **Fabian Groffen (CWI)** Armada, An Evolving Database System
- 2009-19 | **Valentin Robu (CWI)** Modeling Preferences, Strategic Reasoning and Collaboration in Agent-Mediated Electronic Markets
- 2009-20 | **Bob van der Vecht (UU)** Adjustable Autonomy: Controlling Influences on Decision Making
- 2009-21 | **Stijn Vanderlooy (UM)** Ranking and Reliable Classification
- 2009-22 | **Pavel Serdyukov (UT)** Search For Expertise: Going beyond direct evidence
- 2009-23 | **Peter Hofgesang (VU)** Modelling Web Usage in a Changing Environment
- 2009-24 | **Annerieke Heuvelink (VUA)** Cognitive Models for Training Simulations
- 2009-25 | **Alex van Ballegooij (CWI)** "RAM: Array Database Management through Relational Mapping"
- 2009-26 | **Fernando Koch (UU)** An Agent-Based Model for the Development of Intelligent Mobile Services
- 2009-27 | **Christian Glahn (OU)** Contextual Support of social Engagement and Reflection on the Web
- 2009-28 | **Sander Evers (UT)** Sensor Data Management with Probabilistic Models
- 2009-29 | **Stanislav Pokraev (UT)** Model-Driven Semantic Integration of Service-Oriented Applications
- 2009-30 | **Marcin Zukowski (CWI)** Balancing vectorized query execution with bandwidth-optimized storage
- 2009-31 | **Sofiya Katrenko (UVA)** A Closer Look at Learning Relations from Text
- 2009-32 | **Rik Farenhorst (VU) and Remco de Boer (VU)** Architectural Knowledge Management: Supporting Architects and Auditors
- 2009-33 | **Khiet Truong (UT)** How Does Real Affect Affect Affect Recognition In Speech?
- 2009-34 | **Inge van de Weerd (UU)** Advancing in Software Product Management: An Incremental Method Engineering Approach
- 2009-35 | **Wouter Koelewijn (UL)** Privacy en Politiegegevens; Over geautomatiseerde normatieve informatie-uitwisseling
- 2009-36 | **Marco Kalz (OUN)** Placement Support for Learners in Learning Networks
- 2009-37 | **Hendrik Drachslar (OUN)** Navigation Support for Learners in Informal Learning Networks
- 2009-38 | **Riina Vuorikari (OU)** Tags and self-organisation: a metadata ecology for learning resources in a multilingual context
- 2009-39 | **Christian Stahl (TUE, Humboldt-Universitaet zu Berlin)** Service Substitution - A Behavioral Approach Based on Petri Nets
- 2009-40 | **Stephan Raaijmakers (UvT)** Multinomial Language Learning: Investigations into the Geometry of Language

- 2009-41 | **Igor Bereznyy (UvT)** Digital Analysis of Paintings
- 2009-42 | **Toine Bogers (UvT)** Recommender Systems for Social Bookmarking
- 2009-43 | **Virginia Nunes Leal Franqueira (UT)** Finding Multi-step Attacks in Computer Networks using Heuristic Search and Mobile Ambients
- 2009-44 | **Roberto Santana Tapia (UT)** Assessing Business-IT Alignment in Networked Organizations
- 2009-45 | **Jilles Vreeken (UU)** Making Pattern Mining Useful
- 2009-46 | **Loredana Afanasiev (UvA)** Querying XML: Benchmarks and Recursion
- 2010 |**
- 2010-01 | **Matthijs van Leeuwen (UU)** Patterns that Matter
- 2010-02 | **Ingo Wassink (UT)** Work flows in Life Science
- 2010-03 | **Joost Geurts (CWI)** A Document Engineering Model and Processing Framework for Multimedia documents
- 2010-04 | **Olga Kulyk (UT)** Do You Know What I Know? Situational Awareness of Co-located Teams in Multidisplay Environments
- 2010-05 | **Claudia Hauff (UT)** Predicting the Effectiveness of Queries and Retrieval Systems
- 2010-06 | **Sander Bakkes (UvT)** Rapid Adaptation of Video Game AI
- 2010-07 | **Wim Fikkert (UT)** Gesture interaction at a Distance
- 2010-08 | **Krzysztof Siewicz (UL)** Towards an Improved Regulatory Framework of Free Software. Protecting user freedoms in a world of software communities and eGovernments
- 2010-09 | **Hugo Kielman (UL)** A Politiele gegevensverwerking en Privacy, Naar een effectieve waarborging
- 2010-10 | **Rebecca Ong (UL)** Mobile Communication and Protection of Children
- 2010-11 | **Adriaan Ter Mors (TUD)** The world according to MARP: Multi-Agent Route Planning
- 2010-12 | **Susan van den Braak (UU)** Sensemaking software for crime analysis
- 2010-13 | **Gianluigi Folino (RUN)** High Performance Data Mining using Bio-inspired techniques
- 2010-14 | **Sander van Splunter (VU)** Automated Web Service Reconfiguration
- 2010-15 | **Lianne Bodenstaff (UT)** Managing Dependency Relations in Inter-Organizational Models
- 2010-16 | **Sicco Verwer (TUD)** Efficient Identification of Timed Automata, theory and practice
- 2010-17 | **Spyros Kotoulas (VU)** Scalable Discovery of Networked Resources: Algorithms, Infrastructure, Applications
- 2010-18 | **Charlotte Gerritsen (VU)** Caught in the Act: Investigating Crime by Agent-Based Simulation
- 2010-19 | **Henriette Cramer (UvA)** People's Responses to Autonomous and Adaptive Systems
- 2010-20 | **Ivo Swartjes (UT)** Whose Story Is It Anyway? How Improv Informs Agency and Authorship of Emergent Narrative
- 2010-21 | **Harold van Heerde (UT)** Privacy-aware data management by means of data degradation
- 2010-22 | **Michiel Hildebrand (CWI)** End-user Support for Access to Heterogeneous Linked Data
- 2010-23 | **Bas Steunebrink (UU)** The Logical Structure of Emotions
- 2010-24 | **Dmytro Tykhanov ()** Designing Generic and Efficient Negotiation Strategies
- 2010-25 | **Zulfiqar Ali Memon (VU)** Modelling Human-Awareness for Ambient Agents: A Human Mindreading Perspective
- 2010-26 | **Ying Zhang (CWI)** XRPC: Efficient Distributed Query Processing on Heterogeneous XQuery Engines
- 2010-27 | **Marten Voulon (UL)** Automatisch contracteren
- 2010-28 | **Arne Koopman (UU)** Characteristic Relational Patterns
- 2010-29 | **Stratos Idreos(CWI)** Database Cracking: Towards Auto-tuning Database Kernels
- 2010-30 | **Marieke van Erp (UvT)** Accessing Natural History - Discoveries in data cleaning, structuring, and retrieval
- 2010-31 | **Victor de Boer (UVA)** Ontology Enrichment from Heterogeneous Sources on the Web
- 2010-32 | **Marcel Hiel (UvT)** An Adaptive Service Oriented Architecture: Automatically solving Interoperability Problems
- 2010-33 | **Robin Aly (UT)** Modeling Representation Uncertainty in Concept-Based Multimedia Retrieval
- 2010-34 | **Teduh Dirgahayu (UT)** Interaction Design in Service Compositions
- 2010-35 | **Dolf Trieschnigg (UT)** Proof of Concept: Concept-based Biomedical Information Retrieval
- 2010-36 | **Jose Janssen (OU)** Paving the Way for Lifelong Learning: Facilitating competence development through a learning path specification
- 2010-37 | **Niels Lohmann (TUE)** Correctness of services and their composition
- 2010-38 | **Dirk Fahland (TUE)** From Scenarios to components
- 2010-39 | **Ghazanfar Farooq Siddiqui (VU)** Integrative modeling of emotions in virtual agents
- 2010-40 | **Mark van Assem (VU)** Converting and Integrating Vocabularies for the Semantic Web
- 2010-41 | **Guillaume Chaslot (UM)** Monte-Carlo Tree Search
- 2010-42 | **Sybren de Kinderen (VU)** Needs-driven service bundling in a multi-supplier setting - the computational e3-service approach
- 2010-43 | **Peter van Kranenburg (UU)** A Computational Approach to Content-Based Retrieval of Folk Song Melodies
- 2010-44 | **Pieter Bellekens (TUE)** An Approach towards Context-sensitive and User-adapted Access to Heterogeneous Data Sources, Illustrated in the Television Domain
- 2010-45 | **Vasilios Andrikopoulos (UvT)** A theory and model for the evolution of software services
- 2010-46 | **Vincent Pijpers (VU)** e3 alignment: Exploring Inter-Organizational Business-ICT Alignment
- 2010-47 | **Chen Li (UT)** Mining Process Model Variants: Challenges, Techniques, Examples
- 2010-48 | Withdrawn
- 2010-49 | **Jahn-Takeshi Saito (UM)** Solving difficult game positions
- 2010-50 | **Bouke Huurnink (UVA)** Search in Audiovisual Broadcast Archives
- 2010-51 | **Alia Khairia Amin (CWI)** Understanding and supporting information seeking tasks in multiple sources
- 2010-52 | **Peter-Paul van Maanen (VU)** Adaptive Support for Human-Computer Teams: Exploring the Use of Cognitive Models of Trust and Attention
- 2010-53 | **Edgar Meij (UVA)** Combining Concepts and Language Models for Information Access
- 2011 |**
- 2011-01 | **Botond Cseke (RUN)** Variational Algorithms for Bayesian Inference in Latent Gaussian Models
- 2011-02 | **Nick Tinnemeier(UU)** Organizing Agent Organizations. Syntax and Operational Semantics of an Organization-Oriented Programming Language
- 2011-03 | **Jan Martijn van der Werf (TUE)** Compositional Design and Verification of Component-Based Information Systems

- 2011-04 | **Hado van Hasselt (UU)** Insights in Reinforcement Learning: Formal analysis and empirical evaluation of temporal-difference learning algorithms
- 2011-05 | **Base van der Raadt (VU)** Enterprise Architecture Coming of Age - Increasing the Performance of an Emerging Discipline.
- 2011-06 | **Yiwen Wang (TUE)** Semantically-Enhanced Recommendations in Cultural Heritage
- 2011-07 | **Yujia Cao (UT)** Multimodal Information Presentation for High Load Human Computer Interaction
- 2011-08 | **Nieske Vergunst (UU)** BDI-based Generation of Robust Task-Oriented Dialogues
- 2011-09 | **Tim de Jong (OU)** Contextualised Mobile Media for Learning
- 2011-10 | **Bart Bogaert (UvT)** Cloud Content Contention
- 2011-11 | **Dhaval Vyas (UT)** Designing for Awareness: An Experience-focused HCI Perspective
- 2011-12 | **Carmen Bratosin (TUE)** Grid Architecture for Distributed Process Mining
- 2011-13 | **Xiaoyu Mao (UvT)** Airport under Control. Multi-agent Scheduling for Airport Ground Handling
- 2011-14 | **Milan Lovric (EUR)** Behavioral Finance and Agent-Based Artificial Markets
- 2011-15 | **Marijn Koolen (UvA)** The Meaning of Structure: the Value of Link Evidence for Information Retrieval
- 2011-16 | **Maarten Schadd (UM)** Selective Search in Games of Different Complexity
- 2011-17 | **Jiyin He (UVA)** Exploring Topic Structure: Coherence, Diversity and Relatedness
- 2011-18 | **Mark Ponsen (UM)** Strategic Decision-Making in complex games
- 2011-19 | **Ellen Rusman (OU)** The Mind 's Eye on Personal Profiles
- 2011-20 | **Qing Gu (VU)** Guiding service-oriented software engineering - A view-based approach
- 2011-21 | **Linda Terlouw (TUD)** Modularization and Specification of Service-Oriented Systems
- 2011-22 | **Junte Zhang (UVA)** System Evaluation of Archival Description and Access
- 2011-23 | **Wouter Weerkamp (UVA)** Finding People and their Utterances in Social Media
- 2011-24 | **Herwin van Welbergen (UT)** Behavior Generation for Interpersonal Coordination with Virtual Humans On Specifying, Scheduling and Realizing Multimodal Virtual Human Behavior
- 2011-25 | **Syed Waqar ul Qounain Jaffry (VU)** Analysis and Validation of Models for Trust Dynamics
- 2011-26 | **Matthijs Aart Pontier (VU)** Virtual Agents for Human Communication - Emotion Regulation and Involvement-Distance Trade-Offs in Embodied Conversational Agents and Robots
- 2011-27 | **Aniel Bhulai (VU)** Dynamic website optimization through autonomous management of design patterns
- 2011-28 | **Rianne Kaptein(UVA)** Effective Focused Retrieval by Exploiting Query Context and Document Structure
- 2011-29 | **Faisal Kamiran (TUE)** Discrimination-aware Classification
- 2011-30 | **Egon van den Broek (UT)** Affective Signal Processing (ASP): Unraveling the mystery of emotions
- 2011-31 | **Ludo Waltman (EUR)** Computational and Game-Theoretic Approaches for Modeling Bounded Rationality
- 2011-32 | **Nees-Jan van Eck (EUR)** Methodological Advances in Bibliometric Mapping of Science
- 2011-33 | **Tom van der Weide (UU)** Arguing to Motivate Decisions
- 2011-34 | **Paolo Turrini (UU)** Strategic Reasoning in Interdependence: Logical and Game-theoretical Investigations
- 2011-35 | **Maaike Harbers (UU)** Explaining Agent Behavior in Virtual Training
- 2011-36 | **Erik van der Spek (UU)** Experiments in serious game design: a cognitive approach
- 2011-37 | **Adriana Burlutiu (RUN)** Machine Learning for Pairwise Data, Applications for Preference Learning and Supervised Network Inference
- 2011-38 | **Nyree Lemmens (UM)** Bee-inspired Distributed Optimization
- 2011-39 | **Joost Westra (UU)** Organizing Adaptation using Agents in Serious Games
- 2011-40 | **Viktor Clerc (VU)** Architectural Knowledge Management in Global Software Development
- 2011-41 | **Luan Ibraimi (UT)** Cryptographically Enforced Distributed Data Access Control
- 2011-42 | **Michal Sindlar (UU)** Explaining Behavior through Mental State Attribution
- 2011-43 | **Henk van der Schuur (UU)** Process Improvement through Software Operation Knowledge
- 2011-44 | **Boris Reuderink (UT)** Robust Brain-Computer Interfaces
- 2011-45 | **Herman Stehouwer (UvT)** Statistical Language Models for Alternative Sequence Selection
- 2011-46 | **Beibei Hu (TUD)** Towards Contextualized Information Delivery: A Rule-based Architecture for the Domain of Mobile Police Work
- 2011-47 | **Azizi Bin Ab Aziz(VU)** Exploring Computational Models for Intelligent Support of Persons with Depression
- 2011-48 | **Mark Ter Maat (UT)** Response Selection and Turn-taking for a Sensitive Artificial Listening Agent
- 2011-49 | **Andreea Niculescu (UT)** Conversational interfaces for task-oriented spoken dialogues: design aspects influencing interaction quality
- 2012 |**
- 2012-01 | **Terry Kakeeto (UvT)** Relationship Marketing for SMEs in Uganda
- 2012-02 | **Muhammad Umair(VU)** Adaptivity, emotion, and Rationality in Human and Ambient Agent Models
- 2012-03 | **Adam Vanya (VU)** Supporting Architecture Evolution by Mining Software Repositories
- 2012-04 | **Jurriaan Souer (UU)** Development of Content Management System-based Web Applications
- 2012-05 | **Marijn Plomp (UU)** Maturing Interorganizational Information Systems
- 2012-06 | **Wolfgang Reinhardt (OU)** Awareness Support for Knowledge Workers in Research Networks
- 2012-07 | **Rianne van Lambalgen (VU)** When the Going Gets Tough: Exploring Agent-based Models of Human Performance under Demanding Conditions
- 2012-08 | **Gerben de Vries (UVA)** Kernel Methods for Vessel Trajectories
- 2012-09 | **Ricardo Neisse (UT)** Trust and Privacy Management Support for Context-Aware Service Platforms
- 2012-10 | **David Smits (TUE)** Towards a Generic Distributed Adaptive Hypermedia Environment
- 2012-11 | **J.C.B. Rantham Prabhakara (TUE)** Process Mining in the Large: Preprocessing, Discovery, and Diagnostics
- 2012-12 | **Kees van der Sluijs (TUE)** Model Driven Design and Data Integration in Semantic Web Information Systems
- 2012-13 | **Suleman Shahid (UvT)** Fun and Face: Exploring non-verbal expressions of emotion during playful interactions
- 2012-14 | **Evgeny Knutov(TUE)** Generic Adaptation Framework for Unifying Adaptive Web-based Systems

- 2012-15 | **Natalie van der Wal (VU)** Social Agents. Agent-Based Modelling of Integrated Internal and Social Dynamics of Cognitive and Affective Processes.
- 2012-16 | **Fiemke Both (VU)** Helping people by understanding them - Ambient Agents supporting task execution and depression treatment
- 2012-17 | **Amal Elgammal (UvT)** Towards a Comprehensive Framework for Business Process Compliance
- 2012-18 | **Eltjo Poort (VU)** Improving Solution Architecting Practices
- 2012-19 | **Helen Schonenberg (TUE)** What's Next? Operational Support for Business Process Execution
- 2012-20 | **Ali Bahramisharif (RUN)** Covert Visual Spatial Attention, a Robust Paradigm for Brain-Computer Interfacing
- 2012-21 | **Roberto Cornacchia (TUD)** Querying Sparse Matrices for Information Retrieval
- 2012-22 | **Thijs Vis (UvT)** Intelligence, politie en veiligheidsdienst: verenigbare grootheden?
- 2012-23 | **Christian Muehl (UT)** Toward Affective Brain-Computer Interfaces: Exploring the Neurophysiology of Affect during Human Media Interaction
- 2012-24 | **Laurens van der Werff (UT)** Evaluation of Noisy Transcripts for Spoken Document Retrieval
- 2012-25 | **Silja Eckartz (UT)** Managing the Business Case Development in Inter-Organizational IT Projects: A Methodology and its Application
- 2012-26 | **Emile de Maat (UVA)** Making Sense of Legal Text
- 2012-27 | **Hayrettin Gurkok (UT)** Mind the Sheep! User Experience Evaluation & Brain-Computer Interface Games
- 2012-28 | **Nancy Pascall (UvT)** Engendering Technology Empowering Women
- 2012-29 | **Almer Tigelaar (UT)** Peer-to-Peer Information Retrieval
- 2012-30 | **Alina Pommeranz (TUD)** Designing Human-Centered Systems for Reflective Decision Making
- 2012-31 | **Emily Bagarakayo (RUN)** A Learning by Construction Approach for Higher Order Cognitive Skills Improvement, Building Capacity and Infrastructure
- 2012-32 | **Wietske Visser (TUD)** Qualitative multi-criteria preference representation and reasoning
- 2012-33 | **Rory Sie (OUN)** Coalitions in Cooperation Networks (COCOON)
- 2012-34 | **Pavol Jancura (RUN)** Evolutionary analysis in PPI networks and applications
- 2012-35 | **Evert Haasdijk (VU)** Never Too Old To Learn - On-line Evolution of Controllers in Swarm- and Modular Robotics
- 2012-36 | **Denis Ssebugwawo (RUN)** Analysis and Evaluation of Collaborative Modeling Processes
- 2012-37 | **Agnes Nakakawa (RUN)** A Collaboration Process for Enterprise Architecture Creation
- 2012-38 | **Selmar Smit (VU)** Parameter Tuning and Scientific Testing in Evolutionary Algorithms
- 2012-39 | **Hassan Fatemi (UT)** Risk-aware design of value and coordination networks
- 2012-40 | **Agus Gunawan (UvT)** Information Access for SMEs in Indonesia
- 2012-41 | **Sebastian Kelle (OU)** Game Design Patterns for Learning
- 2012-42 | **Dominique Verpoorten (OU)** Reflection Amplifiers in self-regulated Learning
- 2012-43 | Withdrawn
- 2012-44 | **Anna Tordai (VU)** On Combining Alignment Techniques
- 2012-45 | **Benedikt Kratz (UvT)** A Model and Language for Business-aware Transactions
- 2012-46 | **Simon Carter (UVA)** Exploration and Exploitation of Multilingual Data for Statistical Machine Translation
- 2012-47 | **Manos Tsagkias (UVA)** Mining Social Media: Tracking Content and Predicting Behavior
- 2012-48 | **Jorn Bakker (TUE)** Handling Abrupt Changes in Evolving Time-series Data
- 2012-49 | **Michael Kaisers (UM)** Learning against Learning - Evolutionary dynamics of reinforcement learning algorithms in strategic interactions
- 2012-0 | **Steven van Kervel (TUD)** Ontology driven Enterprise Information Systems Engineering
- 2012-51 | **Jeroen de Jong (TUE)** Heuristics in Dynamic Sceduling; a practical framework with a case study in elevator dispatching
- 2013 |**
- 2013-01 | **Viorel Milea (EUR)** News Analytics for Financial Decision Support
- 2013-02 | **Erietta Liarou (CWI)** MonetDB/DataCell: Leveraging the Column-store Database Technology for Efficient and Scalable Stream Processing
- 2013-03 | **Szymon Klarman (VU)** Reasoning with Contexts in Description Logics
- 2013-04 | **Chetan Yadati(TUD)** Coordinating autonomous planning and scheduling
- 2013-05 | **Dulce Pumareja (UT)** Groupware Requirements Evolutions Patterns
- 2013-06 | **Romulo Goncalves(CWI)** The Data Cyclotron: Juggling Data and Queries for a Data Warehouse Audience
- 2013-07 | **Giel van Lankveld (UvT)** Quantifying Individual Player Differences
- 2013-08 | **Robbert-Jan Merk(VU)** Making enemies: cognitive modeling for opponent agents in fighter pilot simulators
- 2013-09 | **Fabio Gori (RUN)** Metagenomic Data Analysis: Computational Methods and Applications
- 2013-10 | **Jeewanie Jayasinghe Arachchige(UvT)** A Unified Modeling Framework for Service Design.
- 2013-11 | **Evangelos Pournaras(TUD)** Multi-level Reconfigurable Self-organization in Overlay Services
- 2013-12 | **Marian Razavian(VU)** Knowledge-driven Migration to Services
- 2013-13 | **Mohammad Safiri(UT)** Service Tailoring: User-centric creation of integrated IT-based homecare services to support independent living of elderly
- 2013-14 | **Jafar Tanha (UVA)** Ensemble Approaches to Semi-Supervised Learning Learning
- 2013-15 | **Daniel Hennes (UM)** Multiagent Learning - Dynamic Games and Applications
- 2013-16 | **Eric Kok (UU)** Exploring the practical benefits of argumentation in multi-agent deliberation
- 2013-17 | **Koen Kok (VU)** The PowerMatcher: Smart Coordination for the Smart Electricity Grid
- 2013-18 | **Jeroen Janssens (UvT)** Outlier Selection and One-Class Classification
- 2013-19 | **Renze Steenhuisen (TUD)** Coordinated Multi-Agent Planning and Scheduling
- 2013-20 | **Katja Hofmann (UvA)** Fast and Reliable Online Learning to Rank for Information Retrieval
- 2013-21 | **Sander Wubben (UvT)** Text-to-text generation by monolingual machine translation
- 2013-22 | **Tom Claassen (RUN)** Causal Discovery and Logic
- 2013-23 | **Patricio de Alencar Silva(UvT)** Value Activity Monitoring
- 2013-24 | **Haitham Bou Ammar (UM)** Automated Transfer in Reinforcement Learning
- 2013-25 | **Agnieszka Anna Latoszek-Berendsen (UM)** Intention-based Decision Support. A new way of representing and implementing clinical guidelines in a Decision Support System

- 2013-26 | **Alireza Zarghami (UT)** Architectural Support for Dynamic Homecare Service Provisioning
- 2013-27 | **Mohammad Huq (UT)** Inference-based Framework Managing Data Provenance
- 2013-28 | **Frans van der Sluis (UT)** When Complexity becomes Interesting: An Inquiry into the Information eXperience
- 2013-29 | **Iwan de Kok (UT)** Listening Heads
- 2013-30 | **Joyce Nakatumba (TUE)** Resource-Aware Business Process Management: Analysis and Support
- 2013-31 | **Dinh Khoa Nguyen (UvT)** Blueprint Model and Language for Engineering Cloud Applications
- 2013-32 | **Kamakshi Rajagopal (OON)** Networking For Learning: The role of Networking in a Lifelong Learner's Professional Development
- 2013-33 | **Qi Gao (TUD)** User Modeling and Personalization in the Microblogging Sphere
- 2013-34 | **Kien Tjin-Kam-Jet (UT)** Distributed Deep Web Search
- 2013-35 | **Abdallah El Ali (Uva)** Minimal Mobile Human Computer Interaction Promotor: Prof. dr. L. Hardman (CWI/UVA)
- 2013-36 | **Than Lam Hoang (TUE)** Pattern Mining in Data Streams
- 2013-37 | **Dirk Beyer (OON)** Ambient Learning Displays
- 2013-38 | **Eelco den Heijer (VU)** Autonomous Evolutionary Art
- 2013-39 | **Joop de Jong (TUD)** A Method for Enterprise Ontology based Design of Enterprise Information Systems
- 2013-40 | **Pim Nijssen (UM)** Monte-Carlo Tree Search for Multi-Player Games
- 2013-41 | **Jochem Liem (UVA)** Supporting the Conceptual Modelling of Dynamic Systems: A Knowledge Engineering Perspective on Qualitative Reasoning
- 2013-42 | **Lijon Planken (TUD)** Algorithms for Simple Temporal Reasoning
- 2013-43 | **Marc Bron (UVA)** Exploration and Contextualization through Interaction and Concepts
- 2014 |**
- 2014-01 | **Nicola Barile (UU)** Studies in Learning Monotone Models from Data
- 2014-02 | **Fiona Tullyano (RUN)** Combining System Dynamics with a Domain Modeling Method
- 2014-03 | **Sergio Raul Duarte Torres (UT)** Information Retrieval for Children: Search Behavior and Solutions
- 2014-04 | **Hanna Jochmann-Mannak (UT)** Websites for children: search strategies and interface design - Three studies on children's search performance and evaluation
- 2014-05 | **Jurriaan van Reijssen (UU)** Knowledge Perspectives on Advancing Dynamic Capability
- 2014-06 | **Damian Tamburri (VU)** Supporting Networked Software Development
- 2014-07 | **Arya Adriansyah (TUE)** Aligning Observed and Modeled Behavior
- 2014-08 | **Samur Araujo (TUD)** Data Integration over Distributed and Heterogeneous Data Endpoints
- 2014-09 | **Philip Jackson (UvT)** Toward Human-Level Artificial Intelligence: Representation and Computation of Meaning in Natural Language
- 2014-10 | **Ivan Salvador Razo Zapata (VU)** Service Value Networks
- 2014-11 | **Janneke van der Zwaan (TUD)** An Empathic Virtual Buddy for Social Support
- 2014-12 | **Willem van Willigen (VU)** Look Ma, No Hands: Aspects of Autonomous Vehicle Control
- 2014-13 | **Arlette van Wissen (VU)** Agent-Based Support for Behavior Change: Models and Applications in Health and Safety Domains
- 2014-14 | **Yangyang Shi (TUD)** Language Models With Meta-information
- 2014-15 | **Natalya Mogles (VU)** Agent-Based Analysis and Support of Human Functioning in Complex Socio-Technical Systems: Applications in Safety and Healthcare
- 2014-16 | **Krystyna Milian (VU)** Supporting trial recruitment and design by automatically interpreting eligibility criteria
- 2014-17 | **Kathrin Dentler (VU)** Computing healthcare quality indicators automatically: Secondary Use of Patient Data and Semantic Interoperability
- 2014-18 | **Mattijs Ghijsen (VU)** Methods and Models for the Design and Study of Dynamic Agent Organizations
- 2014-19 | **Vinicius Ramos (TUE)** Adaptive Hypermedia Courses: Qualitative and Quantitative Evaluation and Tool Support
- 2014-20 | **Mena Habib (UT)** Named Entity Extraction and Disambiguation for Informal Text: The Missing Link
- 2014-21 | **Kassidy Clark (TUD)** Negotiation and Monitoring in Open Environments
- 2014-22 | **Marieke Peeters (UU)** Personalized Educational Games - Developing agent-supported scenario-based training
- 2014-23 | **Eleftherios Sidirouros (UvA/CWI)** Space Efficient Indexes for the Big Data Era
- 2014-24 | **Davide Coelin (VU)** Trusting Semi-structured Web Data
- 2014-25 | **Martijn Lappenschaar (RUN)** New network models for the analysis of disease interaction
- 2014-26 | **Tim Baarslag (TUD)** What to Bid and When to Stop
- 2014-27 | **Rui Jorge Almeida (EUR)** Conditional Density Models Integrating Fuzzy and Probabilistic Representations of Uncertainty
- 2014-28 | **Anna Chmielowiec (VU)** Decentralized k-Clique Matching
- 2014-29 | **Jaap Kabbedijk (UU)** Variability in Multi-Tenant Enterprise Software
- 2014-30 | **Peter de Cock (UvT)** Anticipating Criminal Behaviour
- 2014-31 | **Leo van Moergestel (UU)** Agent Technology in Agile Multiparallel Manufacturing and Product Support
- 2014-32 | **Naser Ayat (UvA)** On Entity Resolution in Probabilistic Data
- 2014-33 | **Tesfa Tegegne (RUN)** Service Discovery in eHealth
- 2014-34 | **Christina Manteli (VU)** The Effect of Governance in Global Software Development: Analyzing Transactive Memory Systems.
- 2014-35 | **Joost van Ooijen (UU)** Cognitive Agents in Virtual Worlds: A Middleware Design Approach
- 2014-36 | **Joos Buijs (TUE)** Flexible Evolutionary Algorithms for Mining Structured Process Models
- 2014-37 | **Maral Dadvar (UT)** Experts and Machines United Against Cyberbullying
- 2014-38 | **Danny Plass-Oude Bos (UT)** Making brain-computer interfaces better: improving usability through post-processing.
- 2014-39 | **Jasmina Maric (UvT)** Web Communities, Immigration, and Social Capital
- 2014-40 | **Walter Omona (RUN)** A Framework for Knowledge Management Using ICT in Higher Education
- 2014-41 | **Frederic Hogenboom (EUR)** Automated Detection of Financial Events in News Text
- 2014-42 | **Carsten Eijkhof (CWI/TUD)** Contextual Multidimensional Relevance Models
- 2014-43 | **Kevin Vlaanderen (UU)** Supporting Process Improvement using Method Increments

2014-44 | **Paulien Meesters (UvT)** Intelligent Blauw. Met als ondertitel: Intelligence-gestuurde politiezorg in gebiedsgebonden eenheden.

2014-45 | **Birgit Schmitz (OUN)** Mobile Games for Learning: A Pattern-Based Approach

2014-46 | **Ke Tao (TUD)** Social Web Data Analytics: Relevance, Redundancy, Diversity

2014-47 | **Shangsong Liang (UVA)** Fusion and Diversification in Information Retrieval

2015 |

2015-01 | **Niels Netten (UvA)** Machine Learning for Relevance of Information in Crisis Response

2015-02 | **Faiza Bukhsh (UvT)** Smart auditing: Innovative Compliance Checking in Customs Controls

2015-03 | **Twan van Laarhoven (RUN)** Machine learning for network data

2015-04 | **Howard Spoelstra (OUN)** Collaborations in Open Learning Environments

2015-05 | **Christoph Bějšch(UT)** Cryptographically Enforced Search Pattern Hiding

2015-06 | **Farideh Heidari (TUD)** Business Process Quality Computation - Computing Non-Functional Requirements to Improve Business Processes

2015-07 | **Maria-Hendrike Peetz(UvA)** Time-Aware Online Reputation Analysis

2015-08 | **Jie Jiang (TUD)** Organizational Compliance: An agent-based model for designing and evaluating organiza-

tional interactions

2015-09 | **Randy Klaassen(UT)** HCI Perspectives on Behavior Change Support Systems

2015-10 | **Henry Hermans (OUN)** OpenU: design of an integrated system to support lifelong learning

2015-11 | **Yongming Luo(TUE)** Designing algorithms for big graph datasets: A study of computing bisimulation and joins

2015-12 | **Julie M. Birkholz (VU)** Modi Operandi of Social Network Dynamics: The Effect of Context on Scientific Collaboration Networks

2015-13 | **Giuseppe Procaccianti(VU)** Energy-Efficient Software

2015-14 | **Bart van Straalen (UT)** A cognitive approach to modeling bad news conversations

2015-15 | **Klaas Andries de Graaf (VU)** Ontology-based Software Architecture Documentation

2015-16 | **Changyun Wei (TUD)** Cognitive Coordination for Cooperative Multi-Robot Teamwork

2015-17 | **André van Cleeff (UT)** Physical and Digital Security Mechanisms: Properties, Combinations and Trade-offs

2015-18 | **Holger Pirk (CWI/UVA)** Waste Not, Want Not! - Managing Relational Data in Asymmetric Memories

2015-19 | **Bernardo Tabuenca (OUN)** Ubiquitous Technology for Lifelong Learners

2015-20 | **Lois Vanhée (UU)** Using Culture and Values to Support Flexible Coordination

Curriculum Vitae

Work Experience

- JUNE 2015-CURRENT | Scientific Programmer at LIP6/THALÈS, Paris
Design of a simulation for modeling multi-sensor systems embedded on large drones. Design of a MAS-based solution for optimizing the use of such a system in a concrete setting
- SEPT 2011-APRIL 2015 | PhD on coordinating artificial agents with artificial cultures UNIVERSITÉ DE MONTPELLIER 2, UTRECHT UNIVERSITY
This thesis is the result of this work period
Promotors: Prof. J. FERBER (Montpellier), Prof. J-J.Ch. MEYER (Utrecht), Co-Promotor: Dr. Frank DIGNUM (Utrecht)
- SEPT 2010-AUG 2011 | Scientific Support at UTRECHT UNIVERSITY, Utrecht
I provided my support to Frank DIGNUM and his students, mostly through the development of software prototypes and writing articles. I also taught Multi-Agent System classes at Utrecht University.
- JUL 2007-AUG 2007 | Technical Support at ANTEA, Orléans
Providing IT support for a large organization (setting-up hardware and software, local and distant support)

Education

- JUL 2010 | Msc. Degree in COMPUTER SCIENCE at ENS CACHAN, Rennes
Tracks: "From Data to Knowledge" and Parallelism
Thesis: "Finding Minimal Winning Coalitions in Turn-Based Games" | Advisor: Pr. S.PINCHINAT, Dr. S.SANNER
- SEPT 2009 | Admitted at ENS CACHAN, Ker Lann
- JUL 2008 | BsC. Degree in COMPUTER SCIENCE at LIFO, Orléans
Specialization: General Informatics
Thesis: "Data Visualization" | Advisor: Dr. M. EXBRAYAT

Teaching

Lecture&Practical	Algorithmic (M1)	Montpellier 2	2012
Practical	Basic Computer Science Tools (L1)	Montpellier 2	2012
Practical	Project Supervision (L2, L3, M1)	Montpellier 2	2012
Practical	Multi-Agent Programming (M1)	Utrecht	2011
Practical	Basic Computer Science Tools (L1)	Montpellier 2	2011
Practical	Agent-Oriented Programming (M1)	Montpellier 2	2011
Practical	Multi-Agent Systems (M1)	Utrecht	2010
Practical	Algorithmic, after-class support (L1)	Orléans	2008