



A Dialogue Game Protocol for Agent Purchase Negotiations

PETER McBURNEY

p.j.mcburney@csc.liv.ac.uk

Department of Computer Science, University of Liverpool, Liverpool L69 7ZF, UK

ROGIER M. VAN EIJK

rogier@cs.uu.nl

Institute of Information and Computing Sciences, Universiteit Utrecht, P.O. Box 80.089, 3508 TB Utrecht, The Netherlands

SIMON PARSONS

sparsonss@mit.edu

Center for Co-ordination Science, Sloan School of Management, Massachusetts Institute of Technology, Cambridge, MA 02142, USA

LEILA AMGOUD

leila.amgoud@irit.fr

Institut de Recherche en Informatique de Toulouse, Université Paul Sabatier, 118 Route de Narbonne, 31062 Toulouse, Cedex 4, France

Abstract. We propose a dialogue game protocol for purchase negotiation dialogues which identifies appropriate speech acts, defines constraints on their utterances, and specifies the different sub-tasks agents need to perform in order to engage in dialogues according to this protocol. Our formalism combines a dialogue game similar to those in the philosophy of argumentation with a model of rational consumer purchase decision behaviour adopted from marketing theory. In addition to the dialogue game protocol, we present a portfolio of decision mechanisms for the participating agents engaged in the dialogue and use these to provide our formalism with an operational semantics. We show that these decision mechanisms are sufficient to generate automated purchase decision dialogues between autonomous software agents interacting according to our proposed dialogue game protocol.

Keywords: argumentation, autonomous agents, consumer decision-making, dialogue games, negotiation.

1. Introduction

Imagine a potential consumer in a developed country seeking to purchase a car. Although she would face a bewildering number of choices as to make and model, she may commence with some idea of the features she wanted. If she is married with a family, the car would need to be large enough to take the whole family. Perhaps, in addition, an estate-car (a station-wagon) would be desirable, in order to carry children, their friends and pets, and their sporting equipment, musical instruments, etc. If the children are young, she may desire safety features, such as child-proof locks on the rear doors. If she will use the car for regular commuting to work, she may require enhanced reliability and comfort. If her husband also plans to drive the car regularly, and he is much taller or much shorter than she is, she may require ready adjustability of the driver's seat, steering-wheel and mirror positions. Her need for such adjustability may be even greater if their teenage children are also to drive the car. If she lives in a country with a hot climate, such as Australia, she may

have strong (and rational) preferences regarding car colour and the presence of front ventilation windows.

With such a list of desired features, it is unlikely that she will know beforehand the extent to which the available makes and models satisfy her requirements. Consequently, her quest for a car will typically take the initial form of a sequence of information searches, aiming to discover what cars are available within her budget, and what specific features they have. Once a certain amount of such information has been collected, we would expect our consumer to evaluate one make and model of car against another. Since it is unlikely that any one specific car will have all the features desired, we would expect her evaluation to involve a number of trade-offs, comparing one less-than-perfect option with another until one is selected.

For most consumers, the list of desired features and their relative importance will change as they collect and process information, and will depend both on the information obtained and its source. For example, many consumers on many purchase occasions are susceptible to what marketers call *word-of-mouth*, opinions from trusted others expressing favour or disfavour for particular features or purchase options.¹ Thus, preferences may change as the result of persuasion by a smooth-talking car-salesman, or from reading an article in a credible magazine. Viewing information-seeking and persuasion as preludes to a transaction leads us to consider the consumer purchase decision process as centered on a sequence of dialogues. These dialogues are undertaken between the potential consumer and other interested parties, such as salespeople and information-brokers.

In this paper, our purpose is to design systems for automated electronic purchase transactions in which such dialogues can occur.² For the purposes of this paper (and following [55]) we define negotiations as interactions concerned with the division of a scarce resource or resources. One approach common in the agent literature is to limit the locutions of the agents involved in a negotiation to quantitative offers and counter-offers — in essence, proposals for the division of the relevant resource — as in [10, 40, 59]. At a minimum, such an approach may involve only three types of locutions: *make_offer(p)*, *accept_offer(p)* and *reject_offer(p)*, where *p* is a quantitative offer. Quantitative locutions such as these can then be generated and assessed automatically on an assumption that each agent seeks to maximize its own utility, as in [9, Chapter 8], or, for agents which are members of the same organization, some combination of its own utility and the estimated utilities of the other participants, as in [11].

Real-life negotiations between human participants, however, are typically richer and more complex than the mere exchange of quantitative offers and counter-offers. Participants request information from each other, collectively seek common information, try to persuade each other of contested propositions, and advance arguments for their own offers and against those of others. This richness has been recognized by the use of argumentation in multi-agent system design, as in [6, 26, 41, 51]; not only may agents present offers in a negotiation, but also the reasons for the offers, any qualifications of and conditions on them, and reactions to them. Recent work has sought to define precisely the protocols specific to such argumentation-based interactions, using dialogue game frameworks [1, 2, 21, 48]. Such protocols allow a participating agent to assert statements in the dialogue and to respond to statements made by other participants. The protocol defines what locutions are possible and the circumstances

under which they may be used. Section 2 discusses dialogue game protocols at a generic level.

While such an argumentation protocol presents an agent with a communications language and the syntax for its use, it does not prescribe when specific locutions should be used by an agent. A dialogue game is therefore not sufficient on its own to generate an automatic discussion between software agents. To do this, we have coupled a dialogue game protocol with a model of consumer purchase decision-making taken from marketing theory. To this consumer purchase model we have added a simple model to generate locutions for the seller-agents engaged in negotiation with consumer-agents. Section 3 describes these models. Section 4 presents our high-level model for purchase negotiation dialogues, for which a dialogue game is specified in Section 5. We discuss the semantics and properties of our formalism in Section 6, and present a worked example in Section 7. The paper concludes with a discussion of related and future work in Section 8.

It is important not to be misled by our use of the word *dialogue*: our focus in this paper is on the design of protocols for automated agent interactions, not on human-machine interactions nor on conversations between humans. A key element of any agent-agent interaction protocol is verifiability: the protocol and its rules must be verifiable on the basis of the actions (including speech-actions) of the participating agents. This implies that the protocol must be defined entirely syntactically, because any semantic element of protocol definitions is never finally verifiable; a sufficiently-clever agent can always insincerely simulate any semantic requirement. Accordingly, we separate syntactic and semantic elements in our definition of the negotiation formalism, with the protocol presented in Section 5 defined entirely in syntactic terms.³ Because of the impossibility of final verification of semantic elements, no participant to a dialogue can know with certainty what another participant really believes. Consequently, an important issue in any dialogue between autonomous agents is inference by participants of each other's beliefs from their statements — and non-statements — in the dialogue. In this paper, we assume participants accept one another's statements at face-value, and leave the question of inference of their true, underlying beliefs for future work.

The formalism we present here concerns just one class of negotiation dialogue, that between potential purchasers and potential sellers of certain categories of consumer products. However, as will be seen, the structure we propose is modular, and so may permit instantiation by different models of purchase-decision and sale. Indeed, different agents engaged in the same negotiation may adopt different purchase-decision or sale-decision models and strategies. Our formal structure may therefore represent a wider class of negotiations, although we do not believe that it could represent all types of negotiation dialogues.

2. Dialogue games

We assume that the agent negotiations occur in some suitable electronic space, which, following [32], we term a *Negotiation Space*.⁴ Argumentation formalisms have focused on the vocabulary and syntax rules for communications between the participants inside such a space. One common approach has adopted the formal dialogue games developed for other

purposes by philosophers of argumentation (e.g. [16, 33]); as in [38], we can summarize the different types of rules of such games at an abstract level as follows:

Commencement Rules: Rules which define the circumstances under which the dialogue begins.

Locution Rules: Rules which specify the nature of utterances permitted in the dialogue. Typically, permitted locutions allow participants to assert propositions, to question or contest other asserted propositions, and to justify previously-asserted propositions. Justifications may involve the presentation of a proof of the proposition or an argument for it, and such presentations may also be legal utterances.⁵ The dialogue game rules may also permit participants to utter propositions to which they assign differing degrees of commitment; for example, one may merely *propose* a proposition, a speech act which entails less commitment than would an *assertion* of the same proposition, as in [37].

Combination Rules: Rules which define the dialogical contexts under which particular locutions are permitted or not, or obligatory or not. For instance, it may not be permitted for a participant to utter the same locution repeatedly, or to assert a proposition and subsequently assert the negation of that proposition in the same dialogue, without in the interim having retracted the former assertion. Similarly, the assertion of a proposition by a participant may oblige that same participant to defend it in defined ways following contestation by other participants.

Commitment Rules: Rules which define the circumstances under which participants express dialogical commitment to a proposition. Typically, the assertion of a claim in the debate is defined as indicating to the other participants some level of commitment to, or support for, the claim, within the context of the dialogue. Establishing a commitment to a claim indicates to the other participants that the speaker will defend that claim against attack within the dialogue, for example, proposing arguments supporting it if requested to do so. In the philosophical tradition of formal dialectics established by Charles Hamblin [16, 17], commitments have no psychological or other meaning outside the dialogue context; in particular, they do not indicate that the speaker necessarily believes the claim. It is standard in this work for dialogue systems to incorporate a public set of commitments, called a *commitment store*, for each participant; these stores are usually non-monotonic, in the sense that participants can also retract committed claims, although usually only under defined circumstances. In our application of automated negotiation, commitments may — if the participants so agree — be taken to indicate a promise to undertake some action outside the dialogue, e.g., a commitment to execute a purchase transaction between the respective parties.

Termination Rules: Rules defining the circumstances under which the dialogue ends.

In Section 5, we present the locutions and associated rules for a dialogue game which implements a purchase negotiation dialogue. However, before doing this, we need to understand the overall structure of such dialogues, and for this we require a model of how potential consumers make purchase decisions and potential sellers make sale decisions. The next section presents models for these decisions, drawing on marketing theory.

3. Marketing models

In this section, we present two models: the first is a model for the consumer's purchase decision, based on the work of marketing theorists over the last four decades in modeling consumer purchase behaviour, [30, Chap. 2] [46]. Such models have been tested empirically and are widely used by marketing managers in industry. The second model we present is a simpler model for the seller's sale decision, which assumes his or her sale decision is rational. Together, these two models will act as high-level specifications for the dialogue game framework to be presented in the next section, by identifying the type of locutions we desire, and the circumstances under which they may be used. Because our model of the consumer purchase decision assumes that a software agent representing a human principal has access to that principal's decision-criteria and preferences, an important issue in the implementation of protocols such as the one we are proposing is the elicitation of the principal's preferences and decision-criteria. This is not necessarily a straightforward task, as the principal may not be able to articulate his or her preferences, particularly for novel products. However, we leave a discussion of this question to another occasion.

3.1. *A consumer decision model*

Economists define commodities as goods for which competing products are distinguishable only on the basis of their price. To marketers, the extent to which potential customers perceive competing products as commodities is evidence of a failure of marketing. Since the work of Kelvin Lancaster [28, 29], marketers have viewed products as bundles of features or attributes, which together form the basis of customer preferences for the product. Thus, to continue the motor vehicle example of Section 1, a car may have attributes such as: maximum speed; acceleration speed; fuel utilization; fuel tank capacity; engine size; passenger safety; child-safety; seating capacity; trunk capacity; the number of doors; anti-theft alarms; a pre-installed phone; air-conditioning; electric windows; colour; price; payment terms; design; after-sales service; warranty period; the availability and costs of spare parts; brand reputation; resale value; etc. Thus both tangible product features (such as engine size) and also intangibles (design, warranty, etc) may be important to a product definition and to consumer preferences for the product.

Different consumers will typically assign different relative importance to these attributes, and consider some to be without relevance for their purchase decision. In addition, even when different consumers assign the same importance to an attribute, they may have different utilities, and hence preferences, for its values, as when two customers think vehicle colour is important, but one customer prefers a red vehicle and the other the identical model in blue. Typically, intending purchasers most prefer feature-bundles which are not available in the marketplace, for instance, desiring a motor vehicle which is very fast, very safe and low-priced. Since their most-preferred bundle is not available, intending purchasers are usually forced to select one, non-optimal, bundle from a set of non-optimal bundles. In these circumstances, the customer purchase decision may be modeled as a multi-attribute trade-off between alternative bundles, no one of which is preferred over all others on all attributes [24]. Market researchers typically use a technique called conjoint

analysis to simulate such complex purchase decision processes, and are thereby able to elicit customer trade-offs between features or feature-bundles [13, 57].

Intending purchasers typically make purchase decisions under conditions of finite information-gathering and information-processing capabilities and often under time constraints.⁶ In general, the time and resource costs of evaluation — what marketers call the *costs of thinking* [50] — mean that a complete evaluative comparison is only justified, if ever, for very important purchases; it is not feasible for most consumers on most purchase occasions. Accordingly, marketing models of consumer purchase behaviour typically assume that full evaluation of competing products is only undertaken on a subset of all those products and brands available for purchase. This set, called the *Consideration Set* or *Evoked Set*, has formed the basis for decision-making models which have been validated empirically, at least in Western marketplaces [18]. As would be expected, consumers typically devote more time and processing effort to those purchase decisions for which they have greater involvement. These are often decisions requiring larger sums of money: for example, more effort may be spent on deciding which make of car to buy than on which brand of perfume or after-shave lotion to buy. However, the monetary value of the transaction is not the only measure of involvement, as for example when a buyer agonizes over the choice of a bottle of perfume or after-shave for a lover.

By definition, not every brand or product makes it into a consumer's consideration set. However, the criteria used for inclusion in the set may not be the same criteria used to evaluate and compare brands once they are inside. Air travelers, for example, may only countenance traveling on airlines with good safety records, but then choose between such airlines on the basis of price or service-friendliness or comfort, etc. One common model for consumer decisions posits criteria for inclusion in the consideration set which are akin to thresholds, as in the airline case; such criteria are called *non-compensatory* because a low score for a product on the attributes specified by the criteria can not be overcome by high scores on other attributes. For example, no amount of price-discounting by a crash-prone airline may be sufficient to induce us to travel on it. Once inside the consideration set, however, evaluation of brands is often assumed by marketers to be undertaken on criteria which are *compensatory*: we trade price for comfort, say, when choosing between safe airlines.

With this understanding of marketing models of consumer behaviour, we adopt the following assumptions for our purchase transaction dialogues. The first four assumptions relate to the space in which purchase negotiations dialogues take place. First, we assume three types of agent roles in our dialogue framework: potential consumer-agents (which we call *buyers* or *consumers*), potential seller-agents (*sellers*) and agents offering information and/or advice to consumers (*advisors*), whether impartially or not.⁷ For ease of reference, we assume each agent has female gender. Second, although category definition is sometimes a difficult practical task in marketing, we assume a single category of products (e.g., motor vehicles) is under consideration by the participants. Extension of our model to more than one category would involve simple indexing of all locutions by category. This would permit simultaneous purchase negotiation dialogues across multiple categories, which may be desirable in some complex purchase situations. Third, we assume that each seller offers one or more products for sale in the Negotiation Space, and that these products can be represented as finite bundles of attributes. Not all bundles may be offered by all sellers. From the seller's viewpoint, the bundles are referred to as *sales-options*; from the consumer's viewpoint, they are referred to as *purchase-options*. Finally, each product

attribute has associated with it a set of values from some finite set. Such values may be quantitative (as in different price-levels) or qualitative (as in the linguistic labels used for the color of a product). For simplicity, we assume that the purchase price of products is uni-dimensional, with values from some non-negative subset of the real numbers, and that this attribute, *price*, is distinguished from the others.

We next adopt five assumptions regarding the nature of the consumer's purchase decision-process. Not all of these assumptions are necessary for the implementation of our system, but they provide useful motivation for the approach we will take in the subsequent sections. First, we assume the purchase decision by a consumer is an individual decision, not a group decision. Thus, each consumer agent acts only for itself or its human principal, and not for or with other agents. Group decision-making would add another level of complexity to the model, an issue we postpone for future work. Second, we assume that the products for which an agent negotiation is being undertaken are high-involvement for the consumer agents concerned. Typically, these will be consumer durables, such as motor vehicles or stereo systems, rather than frequently-purchased goods of low monetary value, such as toothpaste. With this assumption, it will be cost-effective for consumer agents to engage in a purchase negotiation dialogue, devoting time and resources to collecting and rationally evaluating information prior to purchase execution. Low-involvement decisions, by contrast, may be made quickly, on little information or even randomly or whimsically; thus, such decisions may not be amenable to rational argument.⁸

Third, we assume that the consumer purchase decision can be modeled as a two-stage process, in which the first stage is the creation of a consideration set, that is some subset of the products available for purchase (or purchase-options), and where the second stage is an evaluation of those purchase-options in the consideration set. We assume each consumer agent uses one or more non-compensatory criteria for inclusion of competing purchase-options in its consideration set. We call these criteria *inclusion criteria*. For the dialogue modeled in this paper, we will assume that each agent enters the dialogue with such criteria defined and known to itself. We assume each consumer agent uses one or more compensatory criteria for evaluation of those purchase-options included in its consideration set. We call these criteria *selection criteria*; again, we assume that agents enter the dialogue with such criteria defined and known to themselves. For any one agent, inclusion criteria and evaluation criteria will typically differ, and both sets of criteria will differ from one agent to another. The consumer's purchase decision may be based on criteria which are not part of the bundle of attributes of the sales-options as these are presented to the Negotiation Space by the seller agents. For instance, a consumer may not wish to purchase anything from a particular seller agent, due to prior negative experiences with that agent or the perception that the seller has a bad reputation. The sale-option attributes presented by sellers to the Negotiation Space we will refer to as the *displayed* attributes of the option, with other attributes called *non-displayed*.

Fourth, we assume that each consumer agent has a real-valued utility function, which assigns utilities to different purchase-options, for example on the basis of each option's attribute-values. We further assume that the utilities of purchase-options are known to the agent concerned. Agents do not necessarily know the utility functions or valuations of other agents. We assume that all agents are *rational*, in the specific sense of seeking to maximize their perceived expected utilities within their time and information-processing resource constraints. Our final assumption is that consumer agents are able to generate and

assess new options potentially of greater utility than the ones offered to the Negotiation Space by sellers or advisors. One algorithm for this could be as follows:⁹

- 1: Generate all combinations of attributes whose values on the inclusion attributes exceed the threshold levels.
- 2: Of these, consider all those combinations whose values on the selection attributes are greater than or equal to at least one option presented by a seller on at least one selection attribute.
- 3: Calculate the utility of the generated attribute-bundles, and rank them.

We next propose a model for the decision-process of the seller(s) in the transaction, and then, in Section 4, use these models to inform the design of model dialogues between a buyer and a seller.

3.2. *A seller decision model*

We assume that each seller agent only offers bundles to the Negotiation Space that she is willing to supply. Thus, the set of sale-options may differ from one seller to another. For the price attribute, we assume each seller has, for each sale-option, a price-threshold, below which she will not supply the associated product, and that this threshold-value is known by the agent concerned prior to commencement of the negotiation. In general, such thresholds are not public information, and are not revealed explicitly in the course of the negotiation. (Of course, when a seller refuses to supply a product at a particular price, others in the negotiation may be able to infer something about that seller's threshold price-level.) How thresholds are calculated is not important for our model, although a rational agent would be expected to calculate them on some justified basis.¹⁰ Note that the price thresholds will likely differ by bundle, even for the one seller, as for example when a car with optional air-conditioning is sold for more than the identical model without this option.

We further assume that seller-agents have several capabilities regarding the sales-options they offer to potential purchasers:

- They commence the dialogue with at least one pre-determined sales-option.
- During the course of the dialogue, they are empowered to construct and offer new options. Such empowerment could take the form of freedom for the seller agent to offer bundles whose attribute-values are within pre-set ranges, with sale prices for such bundles calculated by pre-determined formulae.
- They are also empowered to assess options presented to the Negotiation Space, either as requests from potential customers or as sales-options from competing sellers, and to evaluate whether such competing options can be matched by new offers from themselves. Again, an agent's autonomy to make such decisions may be limited by pre-set ranges on attribute-values, as with floor price thresholds.

As for consumer agents, we assume that sellers are rational, in the sense of being maximizers of perceived expected utility within time and resource constraints. Consequently, we assume that the sellers are willing to enter into negotiations with any consumer agent willing to purchase from them. Sellers may have selection criteria to be applied to potential

customers before final completion of a transaction, such as customer credit-worthiness or the provision of a cash deposit, analogously to the displayed and non-displayed product attributes used by the purchaser in her decision-making. For this paper, we assume any such criteria used by the seller to select customers are negotiated separately to the main purchase transaction, either before the negotiation modeled here, or afterwards, or both.¹¹

4. A model for purchase negotiation dialogues

We now present high-level models of the dialogues between potential buyers and sellers in a purchase transaction, drawing on the models of purchase decision-making outlined above. We first suppose that a consumer-agent knows, in advance of the transaction, which criteria she will use to form a consideration set and to select items from within this set. Then, given these assumptions, a purchase dialogue could proceed as follows, in a sequence we refer to as *Dialogue 1*, where each stage is labeled as follows.

1. **Open Dialogue:** The dialogue commences.
2. **Inform:** The consumer seeks information from a seller and/or an advisor about what purchase-options are available and their attribute-values. One attribute is assumed to be the price of the purchase option. A seller or an advisor provides such information to the consumer.
3. **Form Consideration Set:** The consumer applies its inclusion criteria to the purchase-options provided and so generates a consideration set. This is achieved (in our model of the consumer) by including those options whose attribute-values are greater than the defined threshold-value on the attributes corresponding to the inclusion criteria. The consumer may continue with this process until all options have been considered, or may cease once the consideration set has reached a certain pre-defined size.¹²
4. **Select Option:** The consumer applies its selection criteria to the purchase-options in the consideration set to generate a preferred option. For our purposes, it does not matter what is the nature of this selection process, provided it generates an ordering of the purchase-options. The consumer agent may, for example, calculate the utility attained by each product on each attribute, and then combine these separate attribute-utilities across each purchase-option, so as to produce an aggregate utility for each purchase-option. Comparison of these aggregate utilities may then generate a rank order of purchase-options.
5. **Negotiate:** Negotiations between the consumer and one or more seller agents are undertaken over the consumer's preferred purchase-options, in order of preference. Given the model we have adopted for the earlier steps in this process, we would expect this stage of negotiation to consist of one or more of the following sequences of interactions:
 - A request by the consumer to purchase a particular option from among those presented from a particular seller.
 - A request by the consumer for an option not thus far presented. For example, a consumer may ask if a product consisting of a novel bundle of attribute-levels is possible, and, if so, what price a seller would seek for this bundle. Sellers may or

may not be willing to entertain such requests, depending on the product category in question, and their own preferences.

- No request by the consumer to purchase a particular option. This may occur, for instance, if no purchase option meets the minimum threshold for inclusion in the consideration set or if all options are priced in excess of the consumer's budget.

Depending on circumstances, more than one of these sequences of interactions may occur.

6. **Confirm:** The participants confirm any purchase agreement they have reached.
7. **Close Dialogue:** The dialogue terminates normally.

Note that stages 3 (**Form Consideration Set**) and 4 (**Select Option**) refer to calculations undertaken by the potential buyer internally; they are not strictly part of a dialogue between potential buyers and sellers. Without these two stages, our model of negotiation dialogue is identical with that proposed by Joris Hulstijn [22], which comprises five stages: (a) Opening the dialogue; (b) Sharing information; (c) Making proposals and counter-proposals; (d) Confirming accepted proposals; (e) Closing the dialogue. We have included the consideration-set formation and the option-selection activities as explicit stages because we wish to extend our model to enable dialogue over these two activities.

Recall that Dialogue 1 assumed that the consumer agent knows in advance its inclusion and selection criteria. This assumption is not realistic, since in many human purchase transactions these criteria sets emerge in the course of the negotiation itself. Suppose, therefore, that the consumer agent seeks (or is instructed by its principal to seek) to establish inclusion and selection criteria in the course of the purchase dialogue. Our model for dialogue would then be as follows, which we call *Dialogue 2*. We denote those stages which are identical with Dialogue 1 with the label *As before*.

1. **Open Dialogue:** *As before:* The dialogue commences.
2. **Inform:** *As before:* The consumer seeks information from a seller and/or an advisor about what purchase-options are available and their attribute-values, and this information is provided.
3. **Seek Criteria:** The consumer seeks information from sellers and/or advisors about what criteria are appropriate for inclusion and evaluation assessments, along with their relative importance weightings, and this information is provided. The consumer may also seek reasons for the suggestions and engage in debate with that agent or agents making the suggestions.
4. **Assess Criteria:** The consumer undertakes a rational assessment of the criteria provided. For our purposes, as with the evaluation of purchase-options against criteria, it does not matter what is the nature of this evaluation process, provided it generates lists of criteria (and thresholds) appropriate for input to the purchase-option inclusion and evaluation assessments.
5. **Form Consideration Set:** *As before:* The consumer applies its inclusion criteria to the purchase-options provided to generate a consideration set.
6. **Select Options:** *As before:* The consumer applies its selection criteria to the purchase-options in the consideration set to generate a preferred option.
7. **Negotiate:** *As before:* Negotiations between the consumer and a seller agent are undertaken over the preferred purchase-option.

8. **Confirm:** *As before:* The participants confirm any purchase agreement they have reached.
9. **Close Dialogue:** *As before:* The dialogue terminates normally.

Automation of Dialogue 2 will require locutions and syntax for argument over preferences and over decision-criteria. That rational arguments are possible between human subjects on such matters is a viewpoint defended cogently in [44], although work will be required to formalize the approaches presented there. We leave that task for another time. We also leave for another time the formal representation of negotiations between buyer agents and advisor agents, and negotiations between seller agents and advisor agents, concerning the desirability, costs or benefits of seeking advice; such subsidiary negotiations could be represented by embedded dialogues, in the manner of [43] or [38]. In this paper, we focus attention on Dialogue 1, and in the next section we present a syntax and semantics to operationalize this model.

Both Dialogue 1 and Dialogue 2 are models of ideal dialogues. We permit participants to enter stages multiple times, in any combination and in any order, subject only to some constraining rules. For Dialogue 1, these rules are as follows:

- The first stage in every purchase dialogue is **Open Dialogue**.
- The **Open Dialogue** stage may occur only once in any purchase dialogue. All other stages may occur more than once.
- The only stages which must occur in every purchase dialogue which terminates normally are **Open Dialogue** and **Close Dialogue**.
- At least one instance of the **Inform** stage must precede the first instance of every other stage, excepting **Open Dialogue** and **Close Dialogue**.
- The **Confirm** stage can only be entered following an instance of the **Negotiate** stage.
- The last stage in every purchase dialogue which terminates normally is the **Close Dialogue** stage.
- Subject only to the constraints expressed in these rules and constraints expressed in the locution-combination rules (articulated below), participants may enter any stage from within any other stage at any time.

We define normal termination of a dialogue in terms of the locutions uttered, which we articulate in the next section. Note that the participants may enter the **Close Dialogue** stage more than once in a particular dialogue. This stage (as the locution-combination rules below will indicate) requires participants to indicate that they wish to leave the dialogue. Thus, this stage remains unconcluded, and the dialogue remains open, whilesoever there are at least one buyer and one seller who wish to continue participating. It is therefore possible for this stage to be entered multiple times in any one dialogue. We now present our proposals for a formal syntax and dialogue game rules for an argumentation game which implements Dialogue 1.

5. A purchase negotiation dialogue game

In this section, we present a dialogue game implementation of Dialogue 1. We list and describe the legal locutions for the dialogue, along with the associated rules for their use.

Our syntax is based on that in [2], modified for the specific consumer purchase domain, and we have been guided in our choices of locutions and combination rules by the principles for rational dialogue between consenting and reasonable participants proposed in [20].

We begin by denoting participating agents by unique identifiers, P_{X1}, P_{X2}, \dots , etc, where $X \in \{B, S, A\}$ denotes the role of the agent as potential buyer, seller or advisor, respectively. As in [40, 51], certain locutions may only be uttered by participants in certain roles. We imagine that an electronic purchase dialogue negotiation space will involve multiple seller and advisor participants who, armed with product options for sale, join any dialogue initiated by a potential buyer participant. Thus, we would expect there to be just one buyer in any dialogue, but most likely many advisors and sellers. However, our design permits multiple numbers of each type of participant. With more than one of each type, the participants may wish to engage in communications with only a subset of the other participants at various times in a dialogue. We enable this by allowing speakers to target their locutions at specific audiences; only those participants specified receive the associated locution. For locutions targeted at all participants, i.e., broadcast communications, we denote the set of all participants by *All*. We next define the elements in the domain of discourse, as follows:

Product Category: We denote product categories by lower-case Greek letters, θ, φ, \dots , and as mentioned in Section 3.1, we assume only one product category is under consideration in any one dialogue. We can imagine that multiple products are potentially available for sale and purchase in any category.

Option Attributes: These are features of products or options, represented by labels of the form a_{iA} . These labels may be finite vectors.

Attribute Values: These are the specific values taken by a particular attribute for a given product or option. The values may be real numbers or elements from some finite set.

Sales and Purchase Options: These are products in a specific product category offered for sale by a seller or advisor agent in a negotiation, or requested for purchase by a buyer or advisor agent. Lower-case letters $\tilde{a}, \tilde{b}, \dots$, etc, early in the Roman alphabet are used for these options. We view these options as bundles of attributes, each attribute taking a specific attribute value. Hence, they are represented by finite vectors of the form: $\tilde{a} = (id, a_{1A}, a_{1B}, a_{2A}, a_{2B}, \dots, a_{nA}, a_{nB})$, where *id* is a unique identifier for the option, and where each a_{iA} is an attribute label and a_{iB} is the value taken by this option on attribute with label a_{iA} . These vectors are the same length for all product-options in a specified category; hence, null values are permitted for those attribute values a_{iB} which are either unknown or not specified.

Propositions: We also assume we have a propositional language, with the usual connectives, whose well-formed formulae are denoted by p, q, r, s, \dots , etc.¹³ These formulae are statements about the other elements in the domain of discourse, such as: "*No red cars are available.*"

As an example of these elements, we may consider a discussion over possible purchase of a car. Here the product category is motor vehicles. One sale or purchase option may be a red Mazda MX3, which has top speed of 140 miles per hour and is offered for sale at \$20,000. This option may be represented by the vector: $(00MX3, colour, red, top_speed,$

140, price, 20000). In this representation, the elements *colour*, *top_speed* and *price* are option attributes. For the first of these attributes, *colour*, the value taken by this option is *red*. Likewise, the value of the attribute *top_speed* is 140, etc.

In the **Inform** stage of the dialogue, participants seek information about what sales options are available. As for the other locutions, the providers of this information may decide which audience they intend to receive it. We assume that each participant P_{Xi} has an Information Store, denoted $IS(P_{Xi})$, which contains the information that P_{Xi} has provided to the dialogue. The entries in the store are 3-tuples (S, P_{Yj}, \tilde{a}) , where S is a set of participants, P_{Yj} is a buyer or seller participant, and \tilde{a} is a sales or purchase option. In the case where P_{Yj} is a seller participant, option \tilde{a} is a sales option which participant P_{Xi} has informed the participants in the set S that seller P_{Yj} is willing to provide. In the case where P_{Yj} is a buyer participant, option \tilde{a} is a purchase option which participant P_{Xi} has informed the participants in the set S that buyer P_{Yj} desires to purchase. This purchase option may not yet be one which a seller has offered to provide. Either participants P_{Xi} and P_{Yj} are identical or, if not, then P_{Xi} is an advisor. Entries are inserted into a particular participant's information stores by a locution uttered by that participant, and the set S is the intended audience for the particular locution. Thus, S indicates the visibility to the dialogue participants of this particular entry in $CS(P_{Xi})$. Entries for which S is the entire set of participants in the dialogue are therefore visible to all participants.

Our model of dialogue has a specific stage, **Confirm**, for participants to confirm agreements they have negotiated. The results of these confirmed agreements are stored in Commitment Stores, denoted $CS(P_{Xi})$ for buyer or seller participant P_{Xi} , and defined similarly to the Information Store. For seller participants, the Commitment Store records commitments made by the seller to sell sales-options to specified buyers. For buyer participants, each store holds commitments made by the buyer to purchase options from specified sellers. As with the Information Stores, the entries in the Commitment Stores are inserted by locutions of the participant concerned, and the visibility of the individual entries depends upon the audience targeted in the locution. Entries in the store $CS(P_{Xi})$ are 3-tuples, (S, P_{Yj}, \tilde{a}) , where S is a set of participants, P_{Yj} is a buyer or seller participant, and \tilde{a} is an option which participant P_{Xi} has committed to sell to or purchase from P_{Yj} , respectively. The set S is the audience to which participant P_{Xi} making the commitment has announced the commitment, and so indicates the visibility to the dialogue participants of this particular entry in $CS(P_{Xi})$.

We now define the locutions in the dialogue game. For each one, we specify any pre-conditions required for its utterance, any responses required and the impacts of the utterance on the information and commitment stores. Note that, for participants embodied with models of the form described in Section 3, each of these locutions will invoke internal responses in the participants. We do not articulate these responses in this section, as they are not strictly part of the interaction protocol. Instead, we articulate these in Section 6, which discusses a semantics for our formalism. Where we denote the speaker of the locution by P_{Xi} , without specifying which type of participant it is, the locution may be uttered by any participant, whatever their role. We group the locutions in terms of the stages of Dialogue 1, although this grouping is somewhat arbitrary, as several locutions could be uttered in more than one stage. The locutions are numbered sequentially, **L1**, **L2**, etc, and we begin with two locutions for the **Open Dialogue** stage.

L1: The **open_dialogue(.)** locution:

Locution: **open_dialogue**(P_{Xi} , *All*, θ), where $X \in \{A, B, S\}$.

Preconditions: This locution must not already have been uttered by a participant within the dialogue. To utter this locution an agent P_{Xi} must have a potential need for a purchase of a product in the specified category, or a willingness to sell or to advise on the sale of products in the category.

Meaning: The speaker, participant P_{Xi} , suggests the opening of a purchase dialogue on product category θ . A dialogue can only commence with this move. The second argument, *All*, indicates that this is a statement broadcast to all participants.

Response: Every other agent P_{Xj} wishing to participate in the dialogue must respond with **enter_dialogue**(P_{Xj} , *All*, θ).

Information Store Updates: No effects.

Commitment Store Updates: No effects.

L2: The **enter_dialogue(.)** locution:

Locution: **enter_dialogue**(P_{Xj} , *All*, θ), where $X \in \{A, B, S\}$.

Preconditions: Within the dialogue, a participant P_{Xi} , with $i \neq j$, must have uttered the locution **open_dialogue**(P_{Xi} , *All*, θ). Participant P_{Xj} must have a potential need for a purchase of a product in the specified category, or a willingness to sell or advise on the sale of products in the category.

Meaning: The speaker, participant P_{Xj} , indicates a willingness to join a purchase negotiation dialogue on product category θ . All intending participants other than the mover of the locution **open_dialogue**(P_{Xi} , *All*, θ) must announce their participation with this move.

Response: No responses required.

Information Store Updates: No effects.

Commitment Store Updates: No effects.

We would expect a typical purchase negotiation dialogue to be initiated by a potential buyer, rather than by sellers or advisors. Note that both these locutions require the speaker to target the utterance to all participants; we thereby preclude the possibility of secret participants in the dialogue.¹⁴ We permit participants to enter the dialogue at any time after the initial locution, but they must declare this entry with an **enter_dialogue(.)** move. Once at least one buyer participant and at least one seller participant have entered the dialogue, we say the dialogue is *open*. Until such time as it is open, the dialogue is said to be *pending*, and no locutions other than **enter_dialogue(.)** and **withdraw_dialogue(.)** are possible. The dialogue remains *open* whosoever there is at least one buyer and at least one seller participating in the dialogue who have not yet uttered a **withdraw_dialogue(.)** locution.

In order to utter any other locutions, the speaker must previously have entered the dialogue. Thus, for all other locutions, there is a precondition that the speaker has previously uttered either the locution **open_dialogue(.)** or the locution **enter_dialogue(.)**. For reasons of space, we do not repeat this general precondition in the following locutions, listing only preconditions specific to the locution. We now present two locutions for the **Inform** stage.

L3: The **seek_info(.)** locution:

Locution: **seek_info**(P_{Xi} , \mathcal{S} , p), for $X \in \{B, A\}$ and \mathcal{S} a set of participants, and p a proposition.

Preconditions: No specific preconditions.

Meaning: The speaker, a consumer or advisor participant P_{Xi} , seeks information from one or more participants in the set \mathcal{S} about what sale-options are available, subject to the constraint expressed by p . For example, the constraint may be a budgetary one, with p expressing the statement that the price is less than some threshold. The constraint may also be a null statement, i.e., expressing no constraints.¹⁵

Response: A seller or advisor participant $P_{Yj} \in \mathcal{S}$ must subsequently utter a **willing_to_sell**(P_{Yj} , \mathcal{T} , P_{Sk} , V) locution, where the elements of the set V of sales options satisfy the constraint p , and where $P_{Xi} \in \mathcal{T}$.

Information Store Updates: No effects.

Commitment Store Updates: No effects.

L4: The **willing_to_sell(.)** locution:

Locution: **willing_to_sell**(P_{Yj} , \mathcal{T} , P_{Sk} , V), for $Y \in \{A, S\}$, \mathcal{T} a set of participants which includes both P_{Yj} and P_{Sk} , where P_{Sk} is a seller participant and V is a set of sales options.

Preconditions: Some participant P_{Xi} must have previously uttered a locution **seek_info**(P_{Xi} , \mathcal{S} , p), where $P_{Yj} \in \mathcal{S}$, and the set of sales options V in the **willing_to_sell(.)** locution must satisfy constraint p .

Meaning: The speaker, a seller or advisor P_{Yj} , indicates to the audience \mathcal{T} a willingness by seller participant P_{Sk} to supply a finite and possibly empty set $V = \{\tilde{a}, \tilde{b}, \dots\}$ of purchase-options to any buyer participant in the set \mathcal{T} . Each of the sales options tendered in the set V must satisfy constraint p uttered as part of the prior **seek_info(.)** locution.

Response: None required.

Information Store Updates: For each $\tilde{a} \in V$, the 3-tuple (\mathcal{T} , P_{Sk} , \tilde{a}) is inserted into $IS(P_{Yj})$, the Information Store for participant P_{Yj} .

Commitment Store Updates: No effects.

The provision of information about sales options by means of the **willing_to_sell(.)** locution does not mean a seller is committed to selling a particular option to a particular buyer. Irrevocable commitment to sale only occurs via the **agree_to_sell(.)** locution, which is presented below. We now present four locutions for the **Negotiate** stage.

L5: The **desire_to_buy(.)** locution:

Locution: **desire_to_buy**(P_{Bi} , \mathcal{S} , \mathcal{T} , V), for P_{Bi} a buyer participant, $\mathcal{T} \subseteq \mathcal{S}$ two sets of participants, and V a set of options.

Preconditions: No specific preconditions. The options included in this utterance need not have been presented in the dialogue before this time.

Meaning: Consumer participant, P_{Bi} , speaking to all the participants in the set \mathcal{S} , requests to purchase an option in the set V of options from any seller in the set \mathcal{T} , where $\mathcal{T} \subseteq \mathcal{S}$.

Response: None required.

Information Store Update: For each $\tilde{a} \in V$ and each $P_{Sk} \in \mathcal{T}$, the 3-tuple $(\mathcal{S}, P_{Sk}, \tilde{a})$ is inserted into $IS(P_{Bi})$, the Information Store for participant P_{Bi} .

Commitment Store Update: No effects.

L6: The **prefer(.)** locution:

Locution: **prefer** $(P_{Bi}, \mathcal{S}, V, W)$, for P_{Bi} a buyer participant, \mathcal{S} a set of participants, and V and W two sets of options.

Preconditions: Each of the sale or purchase options contained in the sets V and W must previously have been included as an option in a **willing_to_sell(.)** locution, for which participant P_{Bi} and every participant in \mathcal{S} was in the intended audience, or a **desire_to_buy(.)** locution, uttered by P_{Bi} to an audience which included \mathcal{S} . Equivalently, we could express this precondition by saying that each of the options contained in V and in W must be elements of an Information Store tuple, a tuple to which P_{Bi} and every participant in \mathcal{S} has viewing access.

Meaning: The speaker, a buyer participant P_{Bi} , indicates to the participants in the set \mathcal{S} that she prefers each option in the finite set V of options to each option in the finite set W .

Response: No response required.

Information Store Update: No effects.

Commitment Store Update: No effects.

L7: The **refuse_to_buy(.)** locution:

Locution: **refuse_to_buy** $(P_{Bi}, \mathcal{S}, \mathcal{T}, W)$, for P_{Bi} a buyer participant, \mathcal{T} a set of seller participants, \mathcal{S} a set of participants such that $\mathcal{T} \subseteq \mathcal{S}$, and W a set of options.

Preconditions: This locution cannot be uttered following a valid utterance of **agree_to_buy** $(P_{Bi}, \mathcal{U}, P_{Sj}, V)$, for which both $P_{Sj} \in \mathcal{T}$ and $V \cap W$ is non-empty.

Meaning: A buyer participant P_{Bi} , speaking to audience \mathcal{S} which includes every participant in the set \mathcal{T} , expresses a refusal to purchase any option in the set W of options from any seller in the set \mathcal{T} of seller participants.

Response: None required.

Information Store Update: No effects.

Commitment Store Update: No effects.

L8: The **refuse_to_sell(.)** locution:

Locution: **refuse_to_sell** $(P_{Sj}, \mathcal{S}, \mathcal{T}, W)$, for P_{Sj} a seller agent, \mathcal{T} a set of buyer participants, \mathcal{S} a set of participants with $\mathcal{T} \subseteq \mathcal{S}$, and W a set of options.

Preconditions: This locution cannot be uttered following a valid utterance of **agree_to_sell** $(P_{Sj}, \mathcal{U}, P_{Bi}, V)$, for which both $P_{Bi} \in \mathcal{T}$ and $V \cap W$ is non-empty.

Meaning: A seller participant P_{Sj} , speaking to audience \mathcal{S} which includes every participant in the set \mathcal{T} , expresses a refusal to sell any option in the set W of options to any buyer in the set \mathcal{T} of buyer participants.

Response: None required.

Information Store Update: No effects.

Commitment Store Update: No effects.

The **prefer(.)** locution enables participants to signal degrees of acceptance of sales-options and purchase-options, thus aiding successful resolution of negotiations. However, as explained in Section 3, we present no mechanism for argument over these preferences in this paper; this will be the subject of future work. The next two locutions express commitments to purchase and sell respectively, and so belong in the **Confirm** stage of the dialogue.

L9: The **agree_to_buy(.)** locution:

Locution: **agree_to_buy**(P_{Bi} , \mathcal{S} , P_{Sj} , V), where P_{Bi} is a buyer participant, \mathcal{S} a set of participants containing P_{Sj} , P_{Sj} is a seller participant, and V a non-empty set of sales options.

Preconditions: For each option $\tilde{a} \in V$, a locution of the form **willing_to_sell**(P_{Yk} , \mathcal{T} , P_{Sj} , W) must previously have been uttered such that $\tilde{a} \in W$, and such that $P_{Bi} \in \mathcal{T}$. In other words, buyer P_{Bi} can only agree to purchase options which have previously been offered to her for sale.

Meaning: Buyer agent P_{Bi} , speaking to audience \mathcal{S} , commits to purchase one of each of the options in the set V from seller agent P_{Sj} . We call P_{Sj} the *intended seller* of the locution.

Response: If seller P_{Sj} is willing to sell some or all of the options in the set V to buyer P_{Bi} , she may respond with an appropriate **agree_to_sell(.)** locution.

Information Store Update: No effects.

Commitment Store Update: For each $\tilde{a} \in V$, the 3-tuple (\mathcal{S} , P_{Sj} , \tilde{a}) is inserted into $CS(P_{Bi})$, the Commitment Store for participant P_{Bi} .

L10: The **agree_to_sell(.)** locution:

Locution: **agree_to_sell**(P_{Sj} , \mathcal{S} , P_{Bi} , V), where P_{Sj} is a seller participant, \mathcal{S} a set of participants containing P_{Bi} , P_{Bi} is a buyer participant, and V a non-empty set of sales options.

Preconditions: For every option $\tilde{a} \in V$, Participant P_{Bi} must previously have uttered the locution **agree_to_buy**(P_{Bi} , \mathcal{S} , P_{Sj} , W) for some set of options W containing \tilde{a} . Note that this condition in turn implies that the options contained in V must previously have been announced to an audience including buyer P_{Bi} through a **willing_to_sell(.)** locution.

Meaning: Seller participant P_{Sj} , speaking to audience \mathcal{S} , commits to selling each of the options contained in the set V to buyer P_{Bi} . We call P_{Bi} the *intended buyer* of the locution.

Response: None required.

Information Store Update: No effect.

Commitment Store Update: For each $\tilde{a} \in V$, the 3-tuple (S, P_{Bi}, \tilde{a}) is inserted into $CS(P_{Sj})$, the Commitment Store for participant P_{Sj} .

Finally, we present a single locution for the **Close Dialogue** stage.

L11: The **withdraw_dialogue(.)** locution:

Locution: **withdraw_dialogue** (P_{Xi}, All, θ) , for $X \in \{A, B, S\}$.

Preconditions: No specific preconditions.

Meaning: The speaker, participant P_{Xi} , announces to all participants her withdrawal from the dialogue negotiating the potential purchase of products in the category θ . This move may be executed at any time following her entry to the dialogue.

Response: None required.

Information Store Update: No effects.

Commitment Store Update: No effects.

In addition to the eleven locutions listed above, it will also be useful in what follows to refer to the *null locution*, which is the act of making no utterance. The purchase negotiation dialogue terminates normally, and the dialogue is said to be *closed*, when that participant withdraws whose departure leaves either no buyer participants or no seller participants remaining in the dialogue. In other words, there must always be at least one buyer and at least one seller participant in a dialogue for it to remain *open*.

We define a commitment to a purchase-transaction as having occurred only after the following sequence of dialogue moves:

agree_to_buy (P_{Bi}, S, P_{Sj}, V)

agree_to_sell (P_{Sj}, S, P_{Bi}, V)

Other locutions, by these or other participants, may be uttered in between these two. Each of these two locutions irrevocably commits the speaker to engage (as buyer or seller, respectively) in a purchase transaction. Because we allow participants to utter **willing_to_sell** and **desire_to_buy** locutions without incurring commitments to engage in a transaction (respectively) to sell or to purchase an option, no specific retraction locution for these two locutions is required; participants may “withdraw a previous statement of a willingness to sell or buy by failing subsequently to execute appropriate **agree_to_sell** or **agree_to_buy** locutions.

Proposition 1 *In the model of Dialogue 1 presented in Section 4, stages 1, 2, 5, 6 and 7 can be executed by judicious choice of these dialogue game locutions.*

Proof: We consider each stage in turn:

1. **Open Dialogue Stage:** This stage commences with an utterance of the locution **open_dialogue** (P_{Xi}, All, θ) and at least one utterance of **enter_dialogue** (P_{Yj}, All, θ) ,

- $j \neq i$. Unless it terminates, the dialogue remains in this stage until at least one buyer and at least one seller enter the dialogue.
2. **Inform Stage:** This stage consists of utterances of **seek_info(.)** and **willing_to_sell(.)** locutions.
 5. **Negotiate Stage:** Negotiation is undertaken through utterances of the locutions, **desire_to_buy(.)**, **prefer(.)**, **refuse_to_buy(.)**, **refuse_to_sell(.)**, along with further use of the **seek_info(.)** and **willing_to_sell(.)** locutions.
 6. **Confirm Stage:** As mentioned above, confirmation of an agreement occurs through use of the two locutions, **agree_to_buy(.)** and **agree_to_sell(.)**, suitably instantiated.
 7. **Close Dialogue Stage:** This stage is entered whenever a participant utters **withdraw_dialogue(.)**. A subsequent utterance of another locution will take the dialogue to a different stage. The **Close Dialogue** stage is only completed when a participant utters **withdraw_dialogue(.)**, and the remaining participants do not include at least both a buyer and a seller participant. \square

The purpose of this Proposition is to show that our proposed dialogue game locutions instantiate the model of a purchase decision negotiation dialogue we presented in Section 4. As we noted in that section, stages 3 and 4 of Dialogue 1 are not strictly stages of the dialogue, since they model calculations which occur inside the buyer participants. Therefore, they are not included in Proposition 1, but are discussed in the next Section.

6. Semantics and automatability

Our definition of the rules of the dialogue game in Section 5 was deliberately exclusively syntactical: we made no assumptions regarding the decision-making architectures or the mental states of the participants before, during or after the dialogue in which they engage. Consequently, any agent willing to submit to the defined rules of the negotiation dialogue may participate in it, regardless of the meaning(s) the agent may place on the locutions uttered. We believe this property ensures wide applicability.¹⁶ In addition, as we mentioned earlier, any semantic element in the dialogue rules is, in any case, never fully verifiable, as a sufficiently clever agent may be able to simulate insincerely any such requirement.

However, one of our objectives is automated dialogues, and here the syntactical rules we have proposed are not sufficient to ensure that dialogues can be generated automatically. To achieve this, we need to vest our individual participants with mechanisms which will invoke particular locutions at particular points in the dialogue, responding to past and anticipated future locutions. We call these mechanisms *semantic decision mechanisms*, although they still may be simulated by the participants, and thus bear little or no relationship to the true decision-making processes or associated “mental states” of the participants. In this section, we first define a portfolio of such internal mechanisms for participating agents which we believe are sufficient for generating automated dialogues, and then, using these, we develop an operational semantics for our dialogue game formalism. Our mechanisms and our semantics draw upon the decision-making models of Section 3. We then consider the formal properties of the framework we have proposed.

6.1. Semantic decision mechanisms

We begin by defining a portfolio of internal decision mechanisms which would enable participating agents to undertake a negotiation dialogue in accordance with the rules of the previous Section. For each mechanism, we first present a high-level functionality, and then discuss implementation of the mechanism. We also identify and label the different outputs of each mechanism, as these will be required for the operational semantics presented subsequently. One possible output under each mechanism is *wait*, which is explained in the discussion of three generic procedures **Do or Wait**, below. Note that the functionalities of the mechanisms for any one agent have some overlap, and so may call upon the same procedures. The mechanisms are grouped according to the type of participants to which they apply: Buyers (**B**), Sellers (**S**) and Advisors (**A**).

B1: Recognize Need: A mechanism which enables the buyer to recognize a need for a purchase in a product category, enabling an agent to initiate or to enter such a dialogue. This mechanism could consist of no more than receipt of an instruction from the agent's human principal, or it could depend on the value of some other parameter, such as inventory levels of current stocks, relative to a pre-determined threshold. We assume three possible outputs for this mechanism: *wait*, *have_need(θ)* and *have_no_need(θ)*, where θ is a product category.

B2: Seek Information: A mechanism to seek information from seller or advisor agents regarding the purchase-options available. This mechanism could be implemented as an automatic request upon commencement of each new negotiation dialogue. This mechanism is assumed to have two outputs: *wait* and *seek_info(θ)*.

B3: Form Consideration Set: A mechanism for a buyer agent to form a consideration set using the information received in the dialogue and the consideration inclusion criteria the agent is assumed to possess. In accordance with the model of Section 3, this mechanism can be implemented by comparing the attribute values of the purchase-options presented in the information received with the threshold-values of the attributes among the inclusion criteria, and then selecting all or some of those options which exceed the thresholds on the designated inclusion attributes. We assume that the selection mechanism is able to deal with null attribute values, for example, by deleting from consideration all those bundles with null values on the attributes relevant to the decision. Note that the attributes referred to in the buyer agent's consideration set-formation process may include both displayed and non-displayed attributes. This mechanism is assumed to have three possible outputs: *wait*, $\mathcal{C}(\theta)$, where $\mathcal{C}(\theta)$ is a non-empty Consideration set of options in the product category θ , and \emptyset , the null set.

B4: Rank Options: A mechanism for a buyer agent to rank a set of purchase options. As we suggest in Section 3.1, such a mechanism could be implemented by defining a real-valued utility function over attribute values of purchase options, and then ordering the options according to their total utility. As with consideration set formation, the attributes used in these calculations could include both displayed and non-displayed attributes. This mechanism is assumed to have two possible outputs: *wait* and $V(\theta)$, where $V(\theta)$ is an ordered set of purchase-options, with the highest-ranked option in the first position.

B5: Select Consideration Set Element: Given a consideration set, and a set of selection criteria, a mechanism for a buyer agent to select one element from the Consideration set. As for mechanism **B4**, this mechanism could be implemented using a real-valued utility function over attribute values of purchase options and then select that element with the greatest utility. If more than one option in the buyer's consideration set receives the same top ranking, we assume the mechanism has some procedure to select one of these, e.g., a random selection. This mechanism is assumed to have two possible outputs: *wait* and $v(\theta)$, where $v(\theta)$ is the selected purchase-option in the product category θ belonging to the buyer's consideration set.

B6: Generate Novel Options: A mechanism to generate novel bundles of attributes, not among those purchase-options already presented to the dialogue by Seller or Advisor Participants. This mechanism could be implemented by constructing new options having greater values on the attributes comprising the inclusion and selection criteria than each of the purchase options already presented. Here, one attribute value would be considered "greater" than another when it results in a higher utility for the option concerned.¹⁷ It may be the case that execution of this mechanism does not generate any novel options, and so we assume three possible outputs for this mechanism: *wait*, the empty set \emptyset and a non-empty set $V(\theta)$ of novel purchase options in the product category θ .

B7: Consider Offers: A mechanism to decide at a particular time whether to: (a) accept at this time one of the purchase options proposed by seller or advisor agents; or (b) reject at this time all the purchase options thus far proposed by seller or advisor agents; or (c) explore at this time potential novel options. Such a mechanism could be implemented by calculating the relative utilities of each purchase option presented, the expected utilities of possible novel options, and the time-dependent utility of a failed negotiation, and choosing that outcome with the highest utility. Computational models for this decision will be similar to those for the **Do or Wait** procedure discussed below. We assume four possible outputs for this mechanism: *wait*, $accept(V(\theta))$, $reject(V(\theta))$ and $explore_novel(\theta)$, where $V(\theta)$ is a set of purchase options in the product category θ .

B8: Consider Withdrawal: A mechanism to enable a Buyer agent to decide to withdraw from the dialogue. To implement such a mechanism, one could incorporate an on-going assessment of the expected utility of negotiating the purchase of an option having utility greater than the time-dependent utility of a failed negotiation. This could be similar to the **Do or Wait** procedure discussed below, although it would need to take into account how many and which other agents were still in the dialogue. We assume this mechanism has two possible outputs: *wait* and $withdraw(\theta)$.

We now present the mechanisms assumed for seller participants. These are, on the whole, much simpler than the Buyer mechanisms, due to the simpler nature of the decision model adopted for seller agents (Section 3.2).

S1: Recognize Category: A mechanism which enables a seller agent to recognize a specific product category as being one of interest. This may mean that the seller currently has products for sale, or that it may simply wish to observe the dialogue which occurs in this category. Thus, the mechanism may be effected by assessing whether the seller has products to sell in the category, and/or whether this is a category

of interest, and/or whether competing sellers or potential buyers are participating. A seller may, for instance, wish to observe all the purchases of an important customer, even when these are in categories outside the seller's own product portfolio. Because we expect a typical dialogue to be initiated by a potential buyer, the seller's mechanism is assumed to be reactive rather than pro-active. We assume three possible outputs for this mechanism: *wait*, *wish_to_enter(θ)* and *wish_not_to_enter(θ)*, where θ is a product category.

S2: Provide Information: A mechanism to provide relevant information concerning available sales options upon receipt of a request from a buyer or advisor agent. This mechanism could be implemented as an automatic response, starting with an initial set of sales options. A seller agent may opt initially only to provide options which are not fully described or not comprehensive of those available for sale, both for reasons of commercial confidentiality and/or because of the nature of the product in question, which may require input from the buyer for its full specification. We assume three possible outputs for this mechanism: *wait*, the empty set \emptyset , and a non-empty set $V(\theta)$ of sales options in the product category θ .

S3: Assess Options: A mechanism for a seller agent to assess whether proposed purchase options presented in a dialogue by a buyer or advisor agent and whether options proposed by competing seller agents can also be offered by the agent. At its simplest, such a mechanism need only comprise a comparison of option attribute values against pre-determined permitted ranges. The permitted ranges for some attributes may depend on the values of other attributes, for example when the sale-price of a car depends on the optional features included in it. The three outputs assumed for this mechanism are *wait*, the empty-set \emptyset and a non-empty set $V(\theta)$ of sales options in the product category θ which the seller agent is able to offer to a potential buyer.

S4: Generate New Options: A mechanism to generate new sales options, on the basis of the permitted values of attributes and on the basis of the competing and proposed options presented to the dialogue. Such a mechanism could consist simply of a rule suggesting every proposed or competing option also be offered, provided the option is assessed as being able to be provided. A similar function is performed in current e-commerce systems by automatic pricebots, which monitor the prices offered by competitors on behalf of a seller and then reset the seller's own prices to be equal or lower than those of the lowest-priced competitor [25]. A more complex mechanism would generate new options dependent on the course of the dialogue. Thus, for example, a seller agent seeking to differentiate its offers from those of competitors [42] may seek to construct new options with attribute values not yet included in options already presented to the dialogue, or with novel combinations of attribute-values. Similarly, a mechanism may generate offers to attract or discourage particular buyer agents; a buyer agent which continually proposes novel options in the one dialogue may not be desirable as a customer, and so a seller agent may determine the set of new options to be offered on the basis of the dialogue history.

We assume this mechanism has two possible outputs: the empty set \emptyset , and a non-empty set $V(\theta)$ of sales options in the product category θ , which the seller agent has not yet offered to the dialogue. Although we have articulated this procedure as a distinct mechanism, our operational semantics for the dialogue game does not invoke it directly; instead, we have assumed that this mechanism is only invoked as a sub-procedure within

the next mechanism, **S5: Decide Offer Tactics**. We distinguish it because of its importance to the execution of the marketing strategy of each seller agent.

S5: Decide Offer Tactics: A mechanism for a seller agent P_{Sj} to decide at a particular time whether to: (a) do nothing; (b) match the options provided by competitor sellers or proposed by buyers; or (c) provide new options, which P_{Sj} has not previously offered. An algorithm to effect this mechanism could run as follows:

1. Undertake an assessment of competitor or buyer options. If P_{Sj} can not offer these, then do nothing. If P_{Sj} can offer these, then proceed to:
2. Attempt to generate new options. If this attempt fails, then either do nothing or offer (some of) the options assessed in the earlier step. If this attempt does not fail (i.e. there are new options which P_{Sj} can offer), then proceed to:
3. Decide to offer (some of) these new options, or offer the same options proposed by competitor sellers or potential customers, or do nothing.

This algorithm is sufficiently generic to incorporate a range of marketing strategies for the seller, e.g., aiming to be a product leader or aiming to match competitors on price, etc. [42, 56]. The algorithm would also permit the marketing strategy to be determined dynamically on the basis of the dialogue history. As the algorithm indicates, this mechanism may invoke some of the other seller mechanisms in its execution. We assume there are three possible outputs to this mechanism: *wait*, the action *do nothing*, and a non-empty set $V(\theta)$ of sales options in the product category θ .

S6: Accept or Reject Offer: A mechanism to decide at a particular time whether to accept or reject an **agree_to_buy(.)** locution made by a buyer agent. This could be implemented by a simple decision rule which indicated acceptance whenever the options proposed by the buyer had values for both displayed and non-displayed attributes falling within the seller's permissible ranges, and rejection otherwise. We assume three outputs for this mechanism: *wait*, $\text{accept}(V(\theta))$ and $\text{reject}(V(\theta))$, where $V(\theta)$ is the set of purchase options in the product category θ indicated in the **agree_to_buy(.)** locution.

S7: Consider Withdrawal: A mechanism to enable a Seller agent to decide to withdraw from the dialogue. This mechanism could be similar to that for Buyer agents, mechanism **B8**. As for Buyers, we assume this mechanism has two possible outputs: *wait* and $\text{withdraw}(\theta)$.

Advisor agents provide advice to buyers, and so the mechanisms they require are a mixture of those required for Buyer and Seller agents, along with an ability to aggregate information they obtain. We therefore omit descriptions of mechanisms identical or nearly so to the ones listed above.

A1: Recognize Category (See mechanism **S1** above.)

A2: Seek Information (See mechanism **B2** above.)

A3: Aggregate Information: A mechanism to aggregate relevant information concerning available sales options. This mechanism could be implemented as a simple concatenation of all information provided by seller agents, or could be edited or summarized in defined ways. Because advisor agents are not assumed to be necessarily impartial, advisors may only include information from sellers satisfying certain criteria. The two possible outputs

of this mechanism are *wait* and a set $V(\theta)$ of purchase options in the product category θ , which one or more sellers are willing to offer to a potential buyer.

A4: Provide Information (See mechanism **S2** above.)

A5: Suggest Novel Options: An advisor may identify an unmet need, based on analysis of the locutions observed in one or more negotiation dialogues.¹⁸ The output of this mechanism is a set $V(\theta)$ of purchase options in the product category θ .

A6: Consider Withdrawal (See mechanism **B8** above.)

In addition to the specific functionalities of the mechanisms listed here, we also assume that each mechanism is equipped with three generic functions. We will discuss these procedures as if they are implemented as components of each of the mechanisms above, although they may just as readily be implemented at some higher, control level.

Do or Wait: A procedure to decide whether or not to initiate the mechanism at this time or to postpone a decision until a future time. In a dialogue, for instance, locutions are uttered at discrete times, and so a buyer agent may form a consideration set before all potential sellers have articulated the sales options they are willing to provide. An overly-hasty buyer may thus incur a potential loss. However, for a buyer to wait too long may also incur costs; a buyer wishing to hire a costume for a fancy-dress party will have no need of the costume if the negotiation does not conclude before the party. Thus, a procedure such as this could be implemented by calculating the expected utility of acting at this time versus that of waiting until a future time, and then choosing either to act or not to act now according to whichever option has the greater expected utility. The rational meta-reasoning architecture of Russell and Wefald [47, Chapter 3] is a model of this kind, and a similar model has been implemented in an agent architecture by Schut and Wooldridge [49]. For this reason, *wait* is an outcome for each of the mechanisms listed above, and this outcome represents an intention by the agent to re-execute the main functionality of the mechanism after a defined, although not necessarily constant, period.

Select Locution: A procedure to decide which, if any, locution to utter, taking as inputs an output state of the mechanism concerned. The possible outputs of this procedure are the valid locutions of Section 5, with the target audience left blank, along with the null locution (i.e remaining silent). Note that *wait* is not an output of this procedure.

Select Target Audience: A procedure to decide the intended target audience for the locution selected by the **Select Locution** procedure. This procedure could be implemented by means of simple rules; for example, the rules for a buyer agent could include: (a) target **seek_info(.)** locutions at all Seller and Advisor participants; (b) target **desire_to_buy(.)** and **prefer(.)** locutions to the largest set of Seller and Advisor participants whose non-display attributes satisfy certain, pre-determined conditions; (c) target **agree_to_buy(.)** locutions at only those Seller and Advisor participants offering the purchase options stated in the locution. Similar rules would apply for Seller and Advisor agents. Such rules, of course, need to comply with the rules of syntax regarding target audiences presented with the locutions. The outputs of this procedure are the valid locutions of Section 5, fully instantiated, along with the null locution. Note, as with the previous procedure, that *wait* is not an output of this procedure.

Although we have presented the functionality of these mechanisms only at a high level, it is clear from the descriptions that, given the consumer-purchase and seller decision models presented in Section 3, they are each readily implementable. Indeed, some of these mechanisms are similar in functionality to those specified for the automated negotiation systems for multi-attribute purchase decisions of [3, 10], although neither system involves argumentation. We discuss this related work in Section 8.

6.2. Operational semantics

We now present an operational semantics [9, 19] for the dialogue game syntax presented in Section 5. An operational semantics indicates how the states of a system change as a result of execution of the commands in a programming language. In our case, the commands in question are the locutions in a negotiation dialogue conducted according the rules of syntax we have presented. We will assume the participating agents are imbued with the semantic mechanisms just described, and the states we will take to be the inputs and outputs of these mechanisms. The locutions uttered in the dialogue effect transitions between states of the mechanisms, as utterances serve as inputs to one or more of the mechanisms of the participating agents, and then these mechanisms in turn produce outputs causing further utterances in the dialogue. Thus, our operational semantics will provide a formal linkage between the dialogue locutions and the semantic mechanisms we have defined, and thus can be used to demonstrate that our protocol can support automated dialogues.¹⁹

To define these links, we allow the ordered 3-tuple $\langle P_{Xi}, \mathbf{K}, s \rangle$ to denote the mechanism with number \mathbf{K} and with an output s of participant P_{Xi} . For ease of presentation, where a transition is invoked by or invokes a particular output of a mechanism \mathbf{K} this is denoted by the specific output s in the third place of the triple; where no specific output is invoked, we denote this by a period in the third place, $\langle P_{Xi}, \mathbf{K}, . \rangle$. Some transitions occur between mechanisms of different agents by means of dialogue locutions; these are denoted by arrows, labeled by the relevant locution number from Section 5. Other transitions occur between the mechanisms of a single agent; these are denoted by unlabeled arrows. We assume in the transition rules below that agents identified in 3-tuples on the right-hand-side of labeled arrows are included in the audience for the particular locution. In the following, we ignore the three generic procedures, **Do or Wait**, **Select Locution** and **Select Target Audience**, associated with each mechanism. Hence, we assume an immediate link between an output of a mechanism and any associated locution, and conversely, without specifying any within-agent transitions involved. Moreover, for simplicity of presentation, we have ignored the advisor agents in this list of transition rules. Because there are no advisor-specific locutions in the dialogue game syntax, extension of our semantics to incorporate advisors is straightforward. As before, we denote the empty set by \emptyset . We define the transition rules as follows, for any product category θ , and for any buyer agent P_{Bi} and any seller agents P_{Sj} and P_{Sk} :

- TR1:** $\langle P_{Bi}, \mathbf{B1}, have_no_need(\theta) \rangle \rightarrow \langle P_{Bi}, \mathbf{B1}, wait \rangle$
TR2: $\langle P_{Bi}, \mathbf{B1}, have_need(\theta) \rangle \xrightarrow{\mathbf{L1}, \mathbf{L2}} \langle P_{Sj}, \mathbf{S1}, . \rangle$
TR3: $\langle P_{Sj}, \mathbf{S1}, wish_not_to_enter(\theta) \rangle \rightarrow \langle P_{Sj}, \mathbf{S1}, wait \rangle$

- TR4:** $\langle P_{S_j}, \mathbf{S1}, wish_to_enter(\theta) \rangle \underline{\mathbf{L2}} \langle P_{B_i}, \mathbf{B2}, \cdot \rangle$
TR5: $\langle P_{B_i}, \mathbf{B2}, seek_info(\theta) \rangle \underline{\mathbf{L3}} \langle P_{S_j}, \mathbf{S2}, \cdot \rangle$
TR6: $\langle P_{S_j}, \mathbf{S2}, \emptyset \rangle \rightarrow \langle P_{S_j}, \mathbf{S2}, wait \rangle$
TR7: $\langle P_{S_j}, \mathbf{S2}, V(\theta) \rangle \underline{\mathbf{L4}} \langle P_{B_i}, \mathbf{B3}, \cdot \rangle, V(\theta) \neq \emptyset.$
TR8: $\langle P_{S_j}, \mathbf{S2}, V(\theta) \rangle \underline{\mathbf{L4}} \langle P_{S_k}, \mathbf{S3}, \cdot \rangle, k \neq j.$
TR9: $\langle P_{S_j}, \mathbf{S3}, \cdot \rangle \rightarrow \langle P_{S_j}, \mathbf{S5}, \cdot \rangle$
TR10: $\langle P_{S_j}, \mathbf{S5}, V(\theta) \rangle \underline{\mathbf{L4}} \langle P_{B_i}, \mathbf{B3}, \cdot \rangle, V(\theta) \neq \emptyset.$
TR11: $\langle P_{S_j}, \mathbf{S5}, do_nothing \rangle \rightarrow \langle P_{S_j}, \mathbf{S7}, \cdot \rangle$
TR12: $\langle P_{B_i}, \mathbf{B3}, \emptyset \rangle \rightarrow \langle P_{B_i}, \mathbf{B7}, \cdot \rangle$
TR13: $\langle P_{B_i}, \mathbf{B7}, accept(V(\theta)) \rangle \underline{\mathbf{L9}} \langle P_{S_j}, \mathbf{S6}, \cdot \rangle, \text{ for } P_{S_j} \text{ the intended seller of } \underline{\mathbf{L9}}.$
TR14: $\langle P_{B_i}, \mathbf{B7}, accept(V(\theta)) \rangle \underline{\mathbf{L9}} \langle P_{S_k}, \mathbf{S7}, \cdot \rangle, \text{ for } P_{S_k} \text{ not the intended seller of } \underline{\mathbf{L9}}.$
TR15: $\langle P_{B_i}, \mathbf{B7}, reject(V(\theta)) \rangle \underline{\mathbf{L7}} \langle P_{S_k}, \mathbf{S5}, \cdot \rangle, \text{ for all seller agents } P_{S_k}.$
TR16: $\langle P_{B_i}, \mathbf{B7}, explore_novel(\theta) \rangle \rightarrow \langle P_{B_i}, \mathbf{B6}, \cdot \rangle$
TR17: $\langle P_{B_i}, \mathbf{B6}, \emptyset \rangle \rightarrow \langle P_{B_i}, \mathbf{B7}, \cdot \rangle$
TR18: $\langle P_{B_i}, \mathbf{B6}, V(\theta) \rangle \underline{\mathbf{L5}} \langle P_{S_k}, \mathbf{S3}, \cdot \rangle$
TR19: $\langle P_{B_i}, \mathbf{B3}, C(\theta) \rangle \rightarrow \langle P_{B_i}, \mathbf{B5}, v(\theta) \rangle, C(\theta) \neq \emptyset.$
TR20: $\langle P_{B_i}, \mathbf{B5}, v(\theta) \rangle \rightarrow \langle P_{B_i}, \mathbf{B7}, \cdot \rangle$
TR21: $\langle P_{S_j}, \mathbf{S6}, accept(V(\theta)) \rangle \underline{\mathbf{L10}} \langle P_{B_i}, \mathbf{B8}, \cdot \rangle, \text{ for } P_{B_i} \text{ the intended buyer of } \underline{\mathbf{L10}}.$
TR22: $\langle P_{S_j}, \mathbf{S6}, accept(V(\theta)) \rangle \underline{\mathbf{L10}} \langle P_{B_k}, \mathbf{B7}, \cdot \rangle, \text{ for } P_{B_k} \text{ not the intended buyer of } \underline{\mathbf{L10}}.$
TR23: $\langle P_{S_j}, \mathbf{S6}, accept(V(\theta)) \rangle \underline{\mathbf{L10}} \langle P_{S_k}, \mathbf{S7}, \cdot \rangle, k \neq j.$
TR24: $\langle P_{S_j}, \mathbf{S6}, reject(V(\theta)) \rangle \underline{\mathbf{L8}} \langle P_{S_k}, \mathbf{S7}, \cdot \rangle, k \neq j.$
TR25: $\langle P_{S_j}, \mathbf{S6}, reject(V(\theta)) \rangle \underline{\mathbf{L8}} \langle P_{B_k}, \mathbf{B7}, \cdot \rangle, \text{ for all buyer agents } P_{B_k}.$
TR26: $\langle P_{B_i}, \mathbf{B7}, reject(V(\theta)) \rangle \rightarrow \langle P_{B_i}, \mathbf{B4}, V(\theta) \rangle$
TR27: $\langle P_{B_i}, \mathbf{B4}, V(\theta) \rangle \underline{\mathbf{L6}} \langle P_{S_j}, \mathbf{S5}, \cdot \rangle$
TR28: $\langle P_{B_i}, \mathbf{B8}, withdraw(\theta) \rangle \underline{\mathbf{L11}} \langle P_{S_j}, \mathbf{S7}, \cdot \rangle$
TR29: $\langle P_{B_i}, \mathbf{B8}, withdraw(\theta) \rangle \underline{\mathbf{L11}} \langle P_{B_l}, \mathbf{B8}, \cdot \rangle, i \neq l.$
TR30: $\langle P_{S_j}, \mathbf{S7}, withdraw(\theta) \rangle \underline{\mathbf{L11}} \langle P_{B_i}, \mathbf{B8}, \cdot \rangle$
TR31: $\langle P_{S_j}, \mathbf{S7}, withdraw(\theta) \rangle \underline{\mathbf{L11}} \langle P_{S_k}, \mathbf{S7}, \cdot \rangle, k \neq j.$
TR32: $\langle P_{S_j}, \mathbf{S4}, \cdot \rangle \rightarrow \langle P_{S_j}, \mathbf{S5}, \cdot \rangle$
TR33: $\langle P_{X_j}, \mathbf{K}, wait \rangle \rightarrow \langle P_{X_j}, \mathbf{K}, \cdot \rangle.$

We offer brief descriptions of these transition rules. Transition Rule **TR1** indicates that a buyer with no need at this time for a product in category θ will not initiate a dialogue, but instead review the situation after some time. Transition Rule **TR2** says that a buyer with a current need for a product in category θ will now initiate a purchase negotiation dialogue by means of locution **L1**, i.e. **open_dialogue(.)**, in the case where such a dialogue is not already initiated, or will enter such a dialogue by means of locution **L2**, i.e., **enter_dialogue(.)**, in the case where it has already been initiated. In either case, the utterance of either of these locutions leads to the execution of mechanism **S1** for each seller agent. Transition Rule **TR3** indicates that a seller which does not wish to enter a dialogue on category θ at this time will wait and review the situation at some point in the future. Rule **TR4** says that a seller which does wish to enter at this time will do so by means of an utterance of locution **L2**, i.e. **enter_dialogue(.)**, and that this utterance will lead each buyer agent to execute mechanism **B2: Seek Information**. When this mechanism leads to an output of *seek_info* in a particular buyer, the buyer is led, according to Transition Rule **TR5**, to utter the locution **L3**, i.e., **seek_info(.)**, an utterance

which in turn invokes mechanism **S2: Provide Information** in each seller already in the dialogue.

Rule **TR6** indicates that a seller agent with no sales options to offer at this time waits and reviews her situation after a suitable time. A seller with sales options $V(\theta)$ at this time, on the other hand, as Transition Rule **TR7** says, utters locution **L4** which indicates to the dialogue a willingness to provide $V(\theta)$. This utterance in turn invokes mechanism **B3: Form Consideration Set** in the participating buyer agents. Transition Rule **TR8** says that utterance of locution **L4** also invokes a mechanism in other seller agents participating in the dialogue, namely mechanism **S3: Assess Options**, in which they compare the set $V(\theta)$ with their own options. Rule **TR9** then indicates that the output of this assessment mechanism becomes an input to the **S5: Decide Offer Tactics** mechanism for the same seller. If this mechanism then results in the seller deciding to offer new options to the dialogue, these are again provided by an execution of locution **L4, willing_to_sell(.)**, as indicated by Transition Rule **TR10**. If the mechanism **S5**, however, decides to do nothing, then, as indicated by Rule **TR11**, the seller agent concerned considers whether or not to withdraw from the dialogue at this time, via mechanism **S7**.

Rule **TR12** states that if the buyer's mechanism **B3: Form Consideration Set** eliminates from consideration all potential purchase options thus far presented to the dialogue, then the buyer will execute mechanism **B7: Consider Offers**. The urgency of the buyer's position may require a purchase transaction even when no purchase satisfies the buyer's inclusion criteria. Transition Rules **TR13** to **TR16** then indicate what occurs as a result of the execution of **B7**. Rule **TR13** indicates that in the case where this mechanism leads to an acceptance by the buyer, then locution **L9, agree_to_buy(.)**, is uttered. For the intended seller specified by buyer, the response to **L9** is an invocation of mechanism **S6: Accept or Reject Offer**. For other sellers, not the intended seller of the locution **L9**, the mechanism invoked is **S7: Consider Withdrawal**, as indicated by Transition Rule **TR14**. Rule **TR15** indicates that a decision by mechanism **B7** to reject purchase options leads, via the locution **L7, refuse_to_buy(.)** to all seller agents considering whether to generate new options, via mechanism **S5**. Rule **TR16** says that a decision by mechanism **B7** to explore novel options invokes mechanism **B6: Generate Novel Options**. Transition Rules **TR17** and **TR18** indicate the effects of this mechanism **B6**. If the output of this mechanism is an empty set, then mechanism **B7: Consider Offers** is invoked, possibly again. If the output of mechanism **B6** is a non-empty set of options, then the buyer concerned utters locution **L5, desire_to_buy(.)**. This utterance results in mechanism **S3: Assess Options** being invoked by all seller agents.

Transition Rule **TR19** returns to mechanism **B3: Form Consideration Set**, in the case where this mechanism results in the formation of a non-empty set. This set then becomes input to mechanism **B5: Select Consideration Set Element**, as shown in **TR17**. Since the Consideration set is assumed to be non-empty and finite, mechanism **B5** will always generate a single option as output. Rule **TR20** then indicates that this output option invokes mechanism **B7: Consider Offers**. The consequences of this invocation have been indicated already, by means of Rules **TR13** through **TR16**.

The next five rules indicate the consequences of a seller invoking mechanism **S6**, which decides whether to accept or reject an offer to purchase from a potential buyer. Rule **TR21** indicates that an acceptance of such an offer leads, via locution **L10, agree_to_sell(.)**, the intended buyer to consider withdrawal from the dialogue, via mechanism **B8**. The buyer

does this because a commitment to purchase has just been executed, and so the dialogue has concluded successfully, at least for this buyer and seller. For other buyers who were not those making the prior **agree_to_buy(.)** locution and thus not the intended buyer of locution **L10**, the successful completion of a transaction is assumed to lead them to consider the offers on the table, via mechanism **B7**, as shown by Transition Rule **TR22**. Likewise, those sellers not involved in this completed transaction are also assumed upon receiving locution **L10** to reconsider the options they have offered to the dialogue and so, as shown in Transition Rule **TR23**, invoke mechanism **S5: Decide Offer Tactics**. The same mechanism is invoked by these other sellers in the case where the intended seller rejects an offer to purchase, as shown in Rule **TR24**. Likewise, the utterance of the locution **L8, refuse_to_sell(.)**, also leads all buyer agents who receive this utterance to consider or reconsider the offers before them, via mechanism **B7**. This is shown by Transition Rule **TR25**.

Transition Rule **TR26** states that when mechanism **B7: Consider Offers** results in a decision by a buyer agent to reject all the purchase options thus far proposed, then the buyer ranks those options currently available, by means of mechanism **B4: Rank Options**. The next Rule, **TR27**, then indicates that such a ranking results in the buyer uttering locution **L6, prefer(.)**, and that this in turn invokes mechanism **S5: Decide Offer Tactics** among those sellers who receive it. These two Transition Rules show that a potential buyer is able to provide sellers with information about her preferences that would not be able to be communicated in a mere exchange of acceptances or rejections of offers.

The four Transition Rules, **TR28** to **TR31**, indicate the effects of decisions to withdraw from the dialogue. Each such decision by an agent leads to an utterance of the locution **L11, withdraw_dialogue(.)**, which in turn leads the remaining participants to consider whether they too should withdraw at this time. Rule **TR32** indicates that the outputs of mechanism **S4: Generate New Options** are always input to mechanism **S5: Decide Offer Tactics**. The final Transition Rule **TR33** states that whenever *wait* is the outcome state of a mechanism **K** of an agent then this results in the same mechanism being executed at a later time, as was stated in the description of the **Do or Wait** procedure.

6.3. Automated dialogues

A primary objective of this research is the design of an argumentation language capable of supporting automated dialogues between autonomous software agents. The design of computational mechanisms to support automated negotiations has also been a recent focus of research in multi-agent systems [9, 10, 23], although this has not employed argumentation mechanisms. In this section, we demonstrate that the dialogue game framework and the semantic mechanisms we have presented are *generative*, i.e., that they can be used by autonomous participating agents to generate dialogues automatically.

Proposition 2 *Autonomous software agents equipped with the functionality of the semantic mechanisms defined in Section 6.1 can engage in automated consumer purchase negotiation dialogues conducted according to the syntactical framework presented in Section 5.*

Proof: Assume we have a set of agents equipped with the semantic mechanisms of Section 6.1, with all notation as before. To prove this result, we need to demonstrate:

(a) that every locution in the dialogue game syntax of Section 5 can be invoked by one or more of the semantic mechanisms of Section 6.1; and (b) that every execution of each of these mechanisms ultimately invokes a locution (which may be the null locution). We show these two results by examining the list of Transition Rules defined in Section 6.2 above. (a) For each locution, we list the mechanisms which invoke them together with (in parentheses) the Transition Rule or Rules which establish this invocation.

- L1:** Mechanism **B1** (Rule **TR2**).
- L2:** Mechanism **B1** (Rule **TR2**); Mechanism **S1** (Rule **TR4**).
- L3:** Mechanism **B2** (Rule **TR5**).
- L4:** Mechanism **S2** (Rules **TR7** and **TR8**); Mechanism **S5** (Rule **TR10**).
- L5:** Mechanism **B6** (Rule **TR18**).
- L6:** Mechanism **B4** (Rule **TR27**).
- L7:** Mechanism **B7** (Rule **TR15**).
- L8:** Mechanism **S6** (Rules **TR24** and **TR25**).
- L9:** Mechanism **B7** (Rules **TR13** and **TR14**).
- L10:** Mechanism **S6** (Rules **TR21**, **TR22** and **TR23**).
- L11:** Mechanism **B8** (Rules **TR28** and **TR29**); Mechanism **S7** (Rules **TR30** and **TR31**).

(b) For each mechanism, we show that every execution either invokes a locution as a direct consequence of the output of the mechanism; or indirectly, by the invocation of another mechanism or mechanisms which ultimately leads to the invocation of a locution. As for part (a), we list the Transition Rules which establish these relationships in parentheses.

- B1: Recognize Need:** Output *have_no_need* invokes mechanism **B1** (Rule **TR1**).
- B1: Recognize Need:** Output *have_need* invokes locutions **L1** and **L2** (Rule **TR2**).
- B2: Seek Information:** Output *seek_info* invokes locution **L3** (Rule **TR5**).
- B3: Form Consideration Set:** Output \emptyset invokes mechanism **B7** (Rule **TR12**).
- B3: Form Consideration Set:** Output $C \neq \emptyset$ invokes mechanism **B5** (Rule **TR19**).
- B4: Rank Options:** Output V invokes locution **L6** (Rule **TR27**).
- B5: Select Consideration Set Element:** Output v invokes mechanism **B7** (Rule **TR20**).
- B6: Generate Novel Options:** Output \emptyset invokes mechanism **B7** (Rule **TR17**).
- B6: Generate Novel Options:** Output $V \neq \emptyset$ invokes locution **L5** (Rule **TR18**).
- B7: Consider Offers:** Output *accept* invokes locution **L9** (Rules **TR13** and **TR14**).
- B7: Consider Offers:** Output *reject* invokes locution **L7** (Rule **TR15**) and mechanism **B4** (Rule **TR26**).
- B7: Consider Offers:** Output *explore_novel* invokes mechanism **B6** (Rule **TR16**).
- B8: Consider Withdrawal:** Output *withdraw* invokes locution **L11** (Rules **TR28** and **TR29**).
- S1: Recognize Category:** Output *wish_not_to_enter* invokes mechanism **S1** (Rule **TR3**).
- S1: Recognize Category:** Output *wish_to_enter* invokes locution **L2** (Rule **TR4**).
- S2: Provide Information:** Output \emptyset invokes mechanism **S2** (Rule **TR6**).
- S2: Provide Information:** Output $V \neq \emptyset$ invokes locution **L4** (Rules **TR7** and **TR8**).
- S3: Assess Options:** Output \emptyset or output $V \neq \emptyset$ invokes mechanism **S5** (Rule **TR9**).
- S4: Generate New Options:** Output \emptyset or output $V \neq \emptyset$ invokes mechanism **S5** (Rule **TR32**).

- S5: Decide Offer Tactics:** Output *do nothing* invokes mechanism **S7** (Rule **TR11**).
- S5: Decide Offer Tactics:** Output $V \neq \emptyset$ invokes locution **L4** (Rule **TR10**).
- S6: Accept or Reject Offer:** Output *accept* invokes locution **L10** (Rules **TR21**, **TR22** and **TR23**).
- S6: Accept or Reject Offer:** Output *reject* invokes locution **L8** (Rules **TR24** and **TR25**).
- S7: Consider Withdrawal:** Output *withdraw* invokes locution **L11** (Rules **TR30** and **TR31**).

A careful examination of this list shows that every mechanism either invokes a locution directly, or invokes a mechanism which invokes a locution, or invokes a mechanism which invokes a further mechanism which invokes a locution, and so on. Note that we have not listed here the *wait* outcome of each mechanism, an outcome which always invokes, after a certain period of time, the same mechanism which generated it. This transition from one mechanism to itself does not invalidate our proof of (b) because it may be seen as invoking the null locution. The entirely silent (empty) dialogue may be viewed as an automated dialogue of null locutions. Similarly, an automated dialogue which is forever silent after some point may be seen as comprising null locutions from this point. Thus, an occurrence of a *wait* outcome by some mechanism, or even an uninterrupted, infinite sequence of such *wait* outcomes, can be seen as generating an automated dialogue consisting of repeated utterances of the null locution. \square

In one sense, this proposition should not be surprising. The consumer purchase decision-model and the seller decision-model are models of how a buyer or seller will act in a purchase negotiation. We have used these models to motivate the design of the syntactical dialogue framework and also for the design of semantic mechanisms for agents participating in such dialogues. Our operational semantics couples these two elements – negotiation dialogue syntax and semantic mechanisms – in a manner consistent with the marketing decision-making models.

7. Example

We now present an annotated example of a consumer purchase negotiation dialogue conducted according to the dialogue game framework we have proposed in Section 5. In this example we do not specify the nature of the product under negotiation, assuming just that it can be described in a logical language. We number the utterances in the dialogue sequence, starting from **U1**. In the annotations to these utterances, we discuss the mechanisms which are invoked by and invoke the locutions, along with transitions between mechanisms themselves. However, we mostly ignore the three generic procedures, **Do or Wait**, **Select Locution** and **Select Target Audience**. In particular, we ignore executions of mechanisms which result in a *wait* outcome from the **Do or Wait** procedure.

We assume that the dialogue comprises three participants, a potential buyer P_{B1} and two potential sellers P_{S1} , P_{S2} . We begin by assuming that P_{B1} executes mechanism **B1: Recognize Need** and that this results in a decision to initiate a consumer purchase negotiation dialogue regarding product category θ . This occurs automatically, according to Transition Rule **TR2**, and invokes an **open_dialogue(.)** locution thus:

U1: open_dialogue(P_{B1}, All, θ)

The receipt of this utterance leads the two seller agents, via Transition Rule **TR2**, to each execute their mechanism **S1: Recognize Category**. By Transition Rule **TR4**, the dialogue then proceeds automatically with the following two locutions:

U2: enter_dialogue(P_{S1}, All, θ)**U3: enter_dialogue**(P_{S2}, All, θ)

By Transition Rule **TR4**, the receipt of these two utterances by P_{B1} invokes mechanism **B2: Seek Information**, which outputs $seek_info(\theta)$. Thus, by **TR5**, buyer P_{B1} next utters **U4**, seeking information on the purchase options available, subject to the constraint expressed by p .

U4: seek_info(P_{B1}, All, p)

Again by Transition Rule **TR5**, receiving this utterance invokes an execution of mechanism **S2: Provide Information** in each seller agent. We assume that each of these agents has a non-empty set of purchase options, respectively $\{\tilde{a}_1, \tilde{a}_2, \tilde{a}_3\}$ and $\{\tilde{b}_1, \tilde{b}_2\}$, which they are willing to provide to a potential buyer and so, by Transition Rule **TR7** they utter the following locutions:

U5: willing_to_sell($P_{S1}, All, P_{S1}, \{\tilde{a}_1, \tilde{a}_2, \tilde{a}_3\}$)**U6: willing_to_sell**($P_{S2}, All, P_{S2}, \{\tilde{b}_1, \tilde{b}_2\}$)

These utterances invoke an execution of mechanism **B3: Form Consideration Set** in the buyer P_{B1} , by Rule **TR7**. We assume the **Do or Wait** procedure of the buyer does not execute the main mechanism until after receipt of both utterances, so that both sets of purchase options are considered by the buyer. We also assume that this process results in the creation of a non-empty consideration set $C_1 = \{\tilde{a}_1, \tilde{a}_3, \tilde{b}_2\}$ by the buyer. By Rule **TR19**, the creation of this set then invokes buyer mechanism **B5: Rank Consideration Set Elements**, producing a single element, say \tilde{a}_3 . By Transition Rule **TR20**, this then invokes buyer mechanism **B7: Consider Offers**. Imagine that execution of this mechanism results in buyer P_{B1} rejecting all the purchase options thus far offered by the two sellers. Transition Rule **TR15** then leads to the buyer's utterance of the following locution:

U7: refuse_to_buy($P_{B1}, All, \{P_{S1}, P_{S2}\}, \{\tilde{a}_1, \tilde{a}_2, \tilde{a}_3, \tilde{b}_1, \tilde{b}_2\}$)

Receipt of this utterance by the two seller agents invokes in each, by Transition Rule **TR15**, the mechanism **S5: Decide Offer Tactics**. Assume that, for seller P_{S1} , this mechanism results in the action *do nothing*, but for seller P_{S2} , the output is a set of two new options, $\{\tilde{b}_3, \tilde{b}_4\}$. For the first seller, this output leads, according to Transition Rule **TR11**, to mechanism **S7: Consider Withdrawal**; we assume the seller decides to stay in the dialogue at this time. For the second seller, the output leads, via Transition Rule **TR10**, to the following utterance:

U8: willing_to_sell($P_{S2}, All, P_{S2}, \{\tilde{b}_3, \tilde{b}_4\}$)

As with the utterances **U5** and **U6**, this utterance leads the buyer to invoke mechanism **B3: Form Consideration Set**. However, the outcome *reject* of buyer mechanism **B7: Consider Offers** prior to utterance **U7** would have also invoked buyer mechanism **B4: Rank Options**, by Transition Rule **TR26**. Therefore, assume that this invocation of mechanism **B3** results in the outcome *wait* while the utterances arising as a consequence of invoking mechanism **B4** are considered. Then, assume that the outcome of the ranking undertaken by **B4** is the ordered set $\{\tilde{a}_3, \tilde{b}_2, \tilde{a}_1, \tilde{a}_2, \tilde{b}_1\}$ (listed from most-preferred to least). Transition Rule **TR27** indicates that

this output then invokes the following utterance by the buyer regarding her preferences:

U9: prefer($P_{B1}, All, \{\tilde{a}_3, \tilde{b}_2\}, \{\tilde{a}_1, \tilde{a}_2, \tilde{b}_1\}$)

As with utterance **U7**, receipt of this utterance by the two seller agents invokes in each, by Transition Rule **TR27**, the mechanism **S5: Decide Offer Tactics**. Assume once again, that for seller P_{S1} , this mechanism results in the action *do nothing*, and that this in turn results in the seller remaining in the dialogue. For seller P_{S2} , suppose that the output of **S5** is the set of new options, $\{\tilde{b}_5, \tilde{b}_6\}$, where $\tilde{b}_5 = \tilde{a}_3$ and \tilde{b}_6 is entirely novel. This output then leads, by Transition Rule **TR10**, to the following utterance:

U10: willing_to_sell($P_{S2}, All, P_{S2}, \{\tilde{b}_5, \tilde{b}_6\}$)

As for utterance **U8**, receipt of this invokes buyer mechanism **B3: Form Consideration Set**. Assume that on this occasion, the outcome of this mechanism is the set $C_2 = \{\tilde{a}_3, \tilde{b}_5, \tilde{b}_6\}$.²⁰ As before, creation of a non-empty consideration set then invokes, by Rule **TR19**, buyer mechanism **B5: Rank Consideration Set Elements**, producing a single element, say \tilde{b}_6 . Transition Rule **TR20** then invokes buyer mechanism **B7: Consider Offers**. Imagine that execution of this mechanism on this occasion results in buyer P_{B1} accepting to purchase \tilde{b}_6 from P_{S2} , and thereby uttering (by Rule **TR13**):

U11: agree_to_buy($P_{B1}, All, P_{S2}, \{\tilde{b}_6\}$)

By the definition of locution **L9**, this locution inserts the triple ($All, P_{S2}, \{\tilde{b}_6\}$) into the Commitment Store of P_{B1} . By Transition Rules **TR13** and **TR14** this invokes mechanism **S6: Accept or Reject Offer** in seller P_{S2} and mechanism **S7: Consider Withdrawal** in seller P_{S1} . Assume seller P_{S1} decides to remain in the dialogue for the moment, so as to observe the reaction of seller P_{S2} .²¹ For P_{S2} , assume the former mechanism leads, by Rule **TR21**, to:

U12: agree_to_sell($P_{S2}, All, P_{B1}, \{\tilde{b}_6\}$)

By the definition of locution **L10**, this locution inserts the triple ($All, P_{B1}, \{\tilde{b}_6\}$) into the Commitment Store of P_{S2} . Moreover, with the exchange of locutions **agree_to_buy** and **agree_to_sell** uttered in **U11** and **U12** between an intended buyer and an intended seller, each speaker has irrevocably committed to a transaction. We can imagine that seller P_{S1} , having received utterance **U12** will now execute mechanism **S7: Consider Withdrawal**, and decide to withdraw from the dialogue. Her last utterance is thus:

U13: withdraw_dialogue(P_{S1}, All, θ)

Similarly, utterance **U12** will have invoked mechanism **B8: Consider Withdrawal** in buyer P_{B1} , by Rule **TR21**. Assume that this leads P_{B1} to utter:

U14: withdraw_dialogue(P_{B1}, All, θ)

By Rules **TR31** and **TR28**, utterances **U13** and **U14** will also have invoked mechanism **S7: Consider Withdrawal** in seller P_{S2} . We assume that, with no one remaining in the dialogue, seller P_{S2} also utters:

U15: withdraw_dialogue(P_{S2}, All, θ)

The dialogue now ends.

Although this example is very simple, it does illustrate some features of the framework we have proposed. Firstly, as the annotation makes clear, the dialogue is completely automated if the agents participating are vested with mechanisms having the functionality we defined in Section 6.1. Secondly, our framework permits what may be considered disorderly dialogues, because some locutions invoke mechanisms simultaneously in multiple participants and

these may have different sequences of consequent invocations. For example, utterance **U7**, where the buyer agent indicates a current refusal to purchase any of the options thus far presented, invokes mechanisms in both the seller agents and in the buyer agent, each of which have consequential invocations. As we note in the annotation following utterance **U8**, these lead to the buyer agent potentially executing both mechanism **B3: Form Consideration Set** and **B4: Rank Options** simultaneously, or nearly so, but on different sets of options. Such multiple executions could potentially lead to many simultaneous and inter-cutting threads in the conversation between the participants.

One solution to this problem would be to enforce a rule analogous to Hitchcock's Orderliness Principle [20], in which only one issue is raised at a time and dealt with before proceeding to others. However, we believe this approach reduces the domain of applicability of the formalism, since the resulting agent dialogues would be much more structured than are human purchase negotiations. Instead, our solution is to allow the participants to exercise judicious and rational use of the **Do or Wait** procedure with every execution of each mechanism. In addition to granting greater autonomy to the participating agents, this has the added advantage of allowing the participants to execute each mechanism or to postpone execution according to an individual and dynamic assessment of their particular costs and benefits at each time in the dialogue.

8. Discussion

This paper has presented a formal dialogue game framework for automated agent dialogues concerning consumer durable purchase negotiations. Our framework is a novel combination of marketing models of consumer and seller decision-making together with a dialogue game framework from the philosophy of argumentation. The use of an argumentation formalism enables richer negotiation dialogues than does a simple exchange of offers and counter-offers. The use of the consumer and seller decision-making models enables these dialogues to be generative, i.e., to be used for automation of dialogues. This has not been a feature of previous agent negotiation models involving argumentation. Thus, the work presented here combines research from philosophy, marketing theory and computer science to produce a novel computational negotiation framework for automated consumer purchase dialogues.

Similar generative mechanisms have been proposed in recent agent negotiation architectures which do not use argumentation. For example, Peyman Faratin [10, Chapter 4] equips agents engaged in automated negotiations with mechanisms for: (a) deciding their responses to multi-attribute offers; (b) proposing new offers involving different trade-offs of the same set of attributes as prior offers; and (c) proposing new offers having different attributes to prior offers. Similarly, Mihai Barbuceanu and Wai-Kau Lo [3] combine a multi-attribute utility model with a constraint optimization solver to enable participants to an automated negotiation to prioritize offers received and to generate responses to them. In [23], such mechanisms for automated negotiation frameworks are called *heuristic* approaches, and are distinguished from approaches using either economic game theory or argumentation. However, the word "heuristic" should not be taken to mean "informal," since the underlying decision models may rest on solid theoretical grounds; the marketing models presented in this paper can be grounded, for example, in maximum-expected-utility decision

theory, as in [45]. Our framework for automated negotiation is the first time such heuristic and argumentation approaches have been combined in a negotiation context.²²

We have assumed throughout this paper that the purchase decision being made is an individual one, a situation which marketing theorists distinguish from purchase decisions made by organizations or groups [5], [30, Ch. 3]. Faratin and colleagues [11] have drawn on marketing models of organizational decision-making to model negotiation over the supply and utilization of multi-attribute services between agents with some interest in common. Because of this shared interest (e.g., the agents may represent different departments within the one company), the formalism is designed to enable agents to seek some form of win-win outcome. Thus, participating agents make trade-offs between attribute levels so as to generate offers with at least the same value to the proposer of the offer but with a higher value to the opponent in the negotiation. In this work, the authors define negotiation as “*a process by which a joint decision is made by two or more parties. The parties first verbalise contradictory demands and then move towards agreements*” [11, p. 119]. In the typology of dialogues proposed by Doug Walton and Erik Krabbe [55], this definition gives the resulting dialogue elements of a deliberation, where the parties share joint responsibility for deciding a course of action, rather than being a pure negotiation, defined as a dialogue where the parties seek to divide some scarce resource.²³

Katia Sycara argues that “*in order to negotiate effectively, agents need the ability to (a) represent and maintain belief models, (b) reason about other agents’ beliefs, and (c) influence other agents’ beliefs and behavior*” [53, p. 204]. Her application domain is that of union-company labor negotiations, a domain significantly less structured than the consumer purchase decisions we are considering. In our case, the agents do not need an explicit ability to reason about each other’s beliefs, but can assume implicitly that each participant in the negotiation desires to maximize its own perceived expected utility (subject to any resource constraints) in the transaction, and will do so through defined interactions. Moreover, Sycara’s system is not capable of automated generation of arguments between autonomous software entities.

Also, as mentioned earlier, recent work by Joris Hulstijn [22] has explored the use of dialogue games as models for dialogues, and we have drawn upon Hulstijn’s five-stage model of negotiation dialogues for our model of a purchase dialogue. Hulstijn [22] also proposes a dialogue game model for information-seeking dialogues, where one participant requests information from another. Other applications of dialogue games to designing systems for agent dialogues have included the work of Frank Dignum and colleagues [6, 7], in which agents seeking to form teams to undertake some joint task engage in persuasion and negotiation dialogues, and the work of several of the present authors and their colleagues treating negotiation dialogues [1, 2], discovery dialogues [36] and deliberation dialogues [21]. Because most human dialogues involve complex mixtures of different types of dialogue (persuasions, negotiations, etc), some attention has also been given to dialogue game models of agent dialogue enabling combinations of dialogue-types [38, 43].

Hulstijn’s motivating example involves a human-machine interaction to purchase theatre tickets, and so research in natural language processing and generation is relevant to that domain. Because, as mentioned throughout, our focus is on artificial dialogues between software entities, this is less relevant to our purpose. However, it is interesting that an influential model of human discourse structure links utterances in a dialogue with the intentions and attentions of participants [14], thereby connecting dialogue locutions

with the mental states of the participants; a computational version of this model has subsequently been proposed [31], enabling it to be used for human-machine interaction. It is possible to view such work as an operational semantics for human language discourses, and thus analogous to the approach we have adopted here.

Similar approaches for inter-agent communications have been proposed recently by Munindar Singh [52] and Frank Guerin and Jeremy Pitt [15], building on speech-act theory. These approaches differ from ours in two main respects. Firstly, they are models of generic dialogues, not specifically purchase negotiations; they draw on typologies of generic locutions from speech act theory which would require specific instantiation to be suitable as protocols for negotiation dialogues between autonomous agents. A key objective of our work is the articulation of such a negotiation-specific protocol. Secondly, participating agents in the models of [15, 52] make public expression of their mental states, for example their beliefs, desires or intentions, relevant to the dialogue. These are called *social commitments*, and using them, locutions in the dialogue can be linked to the mental states of the participants, as in the computational linguistics literature cited in the previous paragraph. Our approach, by contrast, does not require agents to make public expression of their mental states; we therefore preclude the possibility (acknowledged by these authors) that an agent may use such expressions to falsely represent its mental states.

There are several future research directions which we are exploring. First is the further study of the formal properties of the dialogue framework, in particular termination and complexity properties. For example, under what circumstances will dialogues terminate, and in which of these cases will termination be due to the successful negotiation of a transaction? And, can we bound the numbers of locutions needed to reach successful termination? The second area of future research is the extension of the dialogue game framework we have presented to arguments over preferences and decision-criteria, as occurs in Dialogue 2. Formalization of such dialogues, as mentioned above, would be a major undertaking. Thirdly, as explained in Section 6, we have assumed a purchase negotiation dialogue in a specific category commences when a potential buyer recognizes she has a need for a product in that category. We have not modeled the process by which such a need arises. Marketing theoretic models exist for the need-arousal process, and in future work we plan to explore the computational formalization of these. Finally, in this paper we have only considered one purchase decision in isolation, whereas many real-life purchase decisions depend upon or influence other purchases. An example is provided by purchases along manufacturing supply chains, where purchases may be considered in multiple product categories simultaneously. Mechanism design for such multi-level negotiations are explored in [9], motivated by the example of linked purchase of aircraft by airlines and of aircraft-engines by aircraft manufacturers. That work assumed that the participants were seeking to maximize their individual utility by exchanging offers and counter-offers and it included no argumentation component. An obvious extension would thus be the development of a similar dialogue game framework to that presented here.

Acknowledgments

This work was partly funded by the European Union IST Programme, through the Sustainable Lifecycles in Information Ecosystems (SLIE) Project (IST-1999-10948), and by

the British EPSRC, through a PhD studentship. This support is gratefully acknowledged. The first author also wishes to thank John Roberts for introducing him to Consideration Set models in marketing.

Notes

1. Urban et al. [54], for example, develop a model to predict new-car purchase decisions on the basis of word-of-mouth.
2. One could question the extent to which human consumers would be willing to entrust purchase decisions to a software agent, especially for high-involvement purchases. This is part of a larger debate over the rationality of delegation of human decision-making powers to computers [27, 39], a debate we do not enter.
3. In Hitchcock's [20] terminology, the protocol system must satisfy the *externalization* principle, i.e., that all rules are defined in terms of external linguistic behaviour. The dialogue game formalism presented for information-seeking dialogues in [34], for instance, does not satisfy this requirement. For a recent discussion of the problems of semantic verification of agent communication languages, see [58].
4. A similar electronic space for scientific dialogues is called an *Agora* in [35]. Both may be viewed as examples of *Institutions* [40].
5. In some multi-agent system applications of dialogue games, e.g. [1], rationality conditions are imposed on utterances, for example allowing agents to assert statements only when they themselves have a prior argument or proof from their own knowledge base. Such rationality conditions are similar conceptually to the *feasibility pre-condition* in the Agent Communications Language of the Foundation for Intelligent Physical Agents (FIPA) [12], which specifies conditions under which an agent can be considered *sincere* when transmitting a message. For the reasons explained at the end of the previous section, we eschew such conditions.
6. While many observers have argued that electronic markets will reduce search and transaction costs for participants, the additional information provided and its ease of collection may well increase information-processing costs.
7. Note that we are not assuming that advisor agents themselves purchase or sell products, i.e., that they act in a market-making or market-taking capacity, as in [4].
8. In addition, vendors of products which are typically the subject of low-involvement decisions, such as supermarkets, usually do not permit negotiations over the terms of the transaction.
9. Because we assume products are finite bundles of attributes and each attribute takes values only from finite sets, this algorithm would involve only a finite number of steps.
10. For example, on the basis of a utility function which combined assessment of the costs of production and supply of the bundle concerned, expected competitor pricing levels, and/or strategic considerations, such as entry to a new market or the desirability of securing a sale to a particular customer.
11. For example, in the automated fish market designed by Noriega and Sierra [40], each potential purchaser must have a valid credit-status for his or her bid to be accepted by the auctioneer.
12. More complex models are possible. Roberts and Lattin [45], for example, model insertion of a new purchase-option into the consideration set of a consumer on the basis of the difference between the incremental expected benefit of the new option and the estimated additional costs of information search and assessment of it.
13. Note that the use of these symbols for propositional formulae, although standard in the agent argumentation literature, differs from the practice in formal logic, where they often denote atomic variables.
14. We make this assumption for reasons of simplicity. Note that in most consumer marketplaces, sellers are able to observe each other's initial offers, although not always the final deals struck with customers. Moreover, in regulated marketplaces, such as those for telecommunications, sellers usually have to make public filings of their offers.
15. Note that we have only permitted buyer or advisor agents to seek such information. We do this because sellers may be unwilling to provide information to other sellers. While there is no technical reason to stop a seller also participating as a potential buyer through another agent identity, codes of conduct for participation may prevent this happening, at least officially.

16. As an example of a human dialogue which is conducted despite the very different meanings given by the participants to the same locutions, see Friedrich Dürrenmatt's novel, *Die Panne* [8].
17. Of course, using such a mechanism repeatedly in a negotiation dialogue may decrease rather than increase the chances of reaching agreement with sellers.
18. Intermediaries in non-electronic marketplaces often provide value by this means.
19. Note that this linkage does not undermine the purely syntactical definition given to our protocol in Section 5; thus the protocol remains verifiable.
20. Note that we have not assumed that the inclusion criteria for consideration set formation or the criteria for selection of an element from a set are constant through time; it may be that the opportunity or other costs of the time taken to make the decision may alter these criteria. Accordingly, there is no reason to assume that consideration sets constructed at later times by an agent will subsume those constructed earlier.
21. Recall that in Section 5 we required participants to announce their entry to and departure from the dialogue to all participants, so that all are aware of the possibility of being observed by others.
22. In [21], these approaches are also combined in models for deliberation dialogues.
23. To quote Walton and Krabbe: "*The goal of negotiation dialogue is to make a deal. Each participant aims to maximize his share of some goods or services which are in short supply.*" [55, p. 72]

References

1. L. Amgoud, N. Maudet, and S. Parsons, "Modelling dialogues using argumentation," in E. Durfee, (ed.), *Proceedings of the Fourth International Conference on Multi-Agent Systems (ICMAS 2000)*, Boston, MA, USA: IEEE Press, pp. 31–38, 2000.
2. L. Amgoud, S. Parsons, and N. Maudet, "Arguments, dialogue, and negotiation," in W. Horn, (ed.), *Proceedings of the Fourteenth European Conference on Artificial Intelligence (ECAI 2000)*, Berlin, Germany: IOS Press, pp. 338–342, 2000.
3. M. Barbuceanu and W.-K. Lo, "A multi-attribute utility theoretic negotiation architecture for electronic commerce," in C. Sierra, M. Gini, and J. S. Rosenschein, (eds.), *Agents 2000: Proceedings of the Fourth International Conference on Autonomous Agents*, New York, NY, USA: ACM Press, pp. 239–246, 2000.
4. R. A. Bourne and R. Zaidi, "A quote-driven automated market," in M. Schroeder and K. Stathis, (eds.), *Proceedings of the Symposium on Information Agents for E-Commerce, Conference on Artificial Intelligence and the Simulation of Behaviour (AISB-01)*, York, UK: AISB, 2001.
5. K. P. Corfman and S. Gupta, "Mathematical models of group choice and negotiation," in J. Eliashberg and G. L. Lilien, (eds.), *Handbooks in Operations Research and Management Science: Volume 5: Marketing*, North-Holland, Amsterdam, The Netherlands, pp. 83–142, 1993.
6. F. Dignum, B. Dunin-Kępicz, and R. Verbrugge, "Agent theory for team formation by dialogue," in C. Castelfranchi and Y. Lespérance, (eds.), *Pre-Proceedings of the Seventh International Workshop on Agent Theories, Architectures, and Languages (ATAL-2000)*, Boston, USA, pp. 141–156, 2000.
7. F. Dignum, B. Dunin-Kępicz, and R. Verbrugge, "Creating collective intention through dialogue," *Logic Journal of the IGPL*, vol. 9, no. 2, pp. 305–319, 2001.
8. F. Dürrenmatt, *A Dangerous Game*, Jonathan Cape: London, UK, 1960. (Translation by R. and C. Winston from *Die Panne*, published in German in 1956.)
9. R. van Eijk, "Programming languages for agent communications," PhD thesis, Department of Computer Science, Utrecht University: Utrecht, The Netherlands, 2000.
10. P. Faratin, "Automated service negotiation between autonomous computational agents," PhD thesis, Department of Electronic Engineering, Queen Mary and Westfield College, University of London: London, UK, 2000.
11. P. Faratin, C. Sierra, and N. Jennings, "Using similarity criteria to make negotiation trade-offs," in E. Durfee, (ed.), *Proceedings of the Fourth International Conference on Multi-Agent Systems (ICMAS 2000)*, Boston, MA, USA: IEEE Press, pp. 119–126, 2000.
12. FIPA. Communicative Act Library Specification. Technical Report XC00037H, Foundation for Intelligent Physical Agents, 10 August 2001.
13. P. E. Green and A. M. Krieger, "Conjoint analysis with product-positioning applications," in J. Eliashberg and G. L. Lilien, (eds.), *Handbooks in Operations Research and Management Science: Volume 5: Marketing*, North-Holland, Amsterdam, The Netherlands, pp. 467–516, 1993.

14. B. J. Grosz and C. L. Sidner, "Attention, intentions, and the structure of discourse," *Computational Linguistics*, vol. 12, no. 3, pp. 175–204, 1986.
15. F. Guerin and J. Pitt, "Denotational semantics for agent communication languages," in J. P. Müller, E. Andre, S. Sen, and C. Frasson, (eds.), *Proceedings of the Fifth International Conference on Autonomous Agents*, New York, NY, USA: ACM Press, pp. 497–504, 2001.
16. C. L. Hamblin, *Fallacies*, Methuen and Co Ltd: London, UK, 1970.
17. C. L. Hamblin, "Mathematical models of dialogue," *Theoria*, vol. 37, pp. 130–155, 1971.
18. J. R. Hauser and B. Wernerfelt, "An evaluation cost model of consideration sets," *Journal of Consumer Research*, vol. 16, pp. 393–408, 1990.
19. K. V. Hindriks, F. S. de Boer, W. van der Hoek, and J.-J. Ch. Meyer, "Formal semantics for an abstract agent programming language," in M. P. Singh, A. S. Rao, and M. J. Wooldridge, (eds.), *Intelligent Agents IV*, Lecture Notes in Artificial Intelligence 1365, Springer: Berlin, Germany, pp. 215–229, 1998.
20. D. Hitchcock, "Some principles of rational mutual inquiry," in F. van Eemeren, R. Grootendorst, J. A. Blair, and C. A. Willard, (eds.), *Proceedings of the Second International Conference on Argumentation*, Amsterdam, The Netherlands: SICSAT: International Society for the Study of Argumentation, pp. 236–243, 1991.
21. D. Hitchcock, P. McBurney, and S. Parsons, "A framework for deliberation dialogues," in H. V. Hansen, C. W. Tindale, J. A. Blair, and R. H. Johnson, (eds.), *Proceedings of the Fourth Biennial Conference of the Ontario Society for the Study of Argumentation (OSSA 2001)*, Windsor, Ontario, Canada, 2001.
22. J. Hulstijn, "Dialogue models for inquiry and transaction," PhD thesis, Universiteit Twente: The Netherlands, 2000.
23. N. R. Jennings, P. Faratin, A. R. Lomuscio, S. Parsons, M. Wooldridge, and C. Sierra, "Automated negotiation: prospects, methods and challenges," *International Journal of Group Decision and Negotiation*, vol. 10, no. 2, pp. 199–215, 2001.
24. R. L. Keeney and H. Raiffa, *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*, Cambridge University Press: New York City, NY, USA, 1993. First published in 1976.
25. J. O. Kephart and A. R. Greenwald, "Probabilistic pricebots," in S. Parsons and P. Gmytrasiewicz, (eds.), *Game-Theoretic and Decision-Theoretic Agents: Proceedings of the 2001 AAAI Spring Symposium*, AAAI Press: Menlo Park, CA, USA, pp. 37–44, 2001. Technical Report SS-01-03.
26. S. Kraus, K. Sycara, and A. Evenchik, "Reaching agreement through argumentation: a logical model and implementation," *Artificial Intelligence*, vol. 104, pp. 1–69, 1998.
27. A. Kuflik, "Computers in control: rational transfer of authority or irresponsible abdication of autonomy," *Ethics and Information Technology*, vol. 1, no. 3, pp. 173–184, 1999.
28. K. J. Lancaster, "A new approach to consumer theory," *Journal of Political Economy*, vol. 74, pp. 132–157, 1966.
29. K. J. Lancaster, *Consumer Demand: A New Approach*, Columbia University Press: New York City, NY, USA, 1971.
30. G. L. Lilien, P. Kotler, and K. S. Moorthy, *Marketing Models*, Prentice-Hall: Englewood Cliffs, NJ, USA, 1992.
31. K. E. Lochbaum, "A collaborative planning model of intentional structure," *Computational Linguistics*, vol. 24, no. 4, pp. 525–572, 1998.
32. A. R. Lomuscio, M. Wooldridge, and N. R. Jennings, "A classification scheme for negotiation in electronic commerce," in F. Dignum and C. Sierra, (eds.), *Agent-Mediated Electronic Commerce: The European AgentLink Perspective*, Lecture Notes in Computer Science 1991, Springer: Berlin, Germany, pp. 19–33, 2001.
33. J. D. MacKenzie, "Question-begging in non-cumulative systems," *Journal of Philosophical Logic*, vol. 8, pp. 117–133, 1979.
34. N. Maudet and F. Evrard, "A generic framework for dialogue game implementation," in *Proceedings of the Second Workshop on Formal Semantics and Pragmatics of Dialog*, Thirteenth International Twente Workshop on Language Technology (TWLT 13), Universiteit Twente, The Netherlands: Centre for Telematics and Information Technology, 1998.
35. P. McBurney and S. Parsons, "Risk agoras: Dialectical argumentation for scientific reasoning," in C. Boutilier and M. Goldszmidt, (eds.), *Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence (UAI-2000)*, San Francisco, CA, USA: Morgan Kaufmann, pp. 371–379, 2000.
36. P. McBurney and S. Parsons, "Chance discovery using dialectical argumentation," in T. Terano, T. Nishida, A. Namatame, S. Tsumoto, Y. Ohsawa, and T. Washio, (eds.), *New Frontiers in Artificial Intelligence: Joint*

- JSAI 2001 Workshop Post Proceedings*, Lecture Notes in Artificial Intelligence 2253, Springer: Berlin, Germany, pp. 414–424, 2001.
37. P. McBurney and S. Parsons, “Representing epistemic uncertainty by means of dialectical argumentation,” *Annals of Mathematics and Artificial Intelligence*, vol. 32, nos. 1–4, pp. 125–169, 2001.
 38. P. McBurney and S. Parsons, “Games that agents play: A formal framework for dialogues between autonomous agents,” *Journal of Logic, Language and Information*, vol. 11, no. 3, pp. 315–334, 2002.
 39. J. Moor, “Are there decisions computers should never make?” *Nature and System*, vol. 1, pp. 217–229, 1979.
 40. P. Noriega and C. Sierra, “Towards layered dialogical agents,” in J. P. Müller, M. J. Wooldridge, N. R. Jennings, (eds.), *Intelligent Agents III: Agent Theories, Architectures, and Languages*, Lecture Notes in Artificial Intelligence 1193, Springer: Berlin, Germany, pp. 173–188, 1997.
 41. S. Parsons, C. Sierra, and N. R. Jennings, “Agents that reason and negotiate by arguing,” *Journal of Logic and Computation*, vol. 8, no. 3, pp. 261–292, 1998.
 42. M. E. Porter, *Competitive Advantage: Creating and Sustaining Superior Performance*, Free Press: New York, NY, USA, 1985.
 43. C. Reed, “Dialogue frames in agent communications,” in Y. Demazeau, (ed.), *Proceedings of the Third International Conference on Multi-Agent Systems (ICMAS-98)*, IEEE Press, pp. 246–253, 1998.
 44. H. S. Richardson, *Practical Reasoning about Final Ends*, Cambridge University Press: Cambridge, UK, 1994.
 45. J. H. Roberts and J. M. Lattin, “Development and testing of a model of consideration set composition,” *Journal of Marketing Research*, vol. 28, pp. 429–440, 1991.
 46. J. H. Roberts and G. L. Lilien, “Explanatory and predictive models of consumer behavior,” in J. Eliashberg and G. L. Lilien, (eds.), *Handbooks in Operations Research and Management Science: Volume 5: Marketing*, North-Holland, Amsterdam, The Netherlands, pp. 27–82, 1993.
 47. S. Russell and E. Wefald, *Do The Right Thing: Studies in Limited Rationality*, MIT Press: Cambridge, MA, USA, 1991.
 48. F. Sadri, F. Toni, and P. Torroni, “Logic agents, dialogues and negotiation: An abductive approach,” in M. Schroeder and K. Stathis, (eds.), *Proceedings of the Symposium on Information Agents for E-Commerce, Artificial Intelligence and the Simulation of Behaviour Conference (AISB-2001)*, York, UK: AISB, 2001.
 49. M. C. Schut and M. J. Wooldridge, “Intention reconsideration in complex environments,” in C. Sierra, M. Gini, and J. Rosenschein, (eds.), *Agents 2000: Proceedings of the Fourth International Conference on Autonomous Agents*, New York, NY, USA: ACM Press, pp. 209–216, 2000.
 50. S. M. Shugan, “The cost of thinking,” *Journal of Consumer Research*, vol. 7, pp. 99–111, 1980.
 51. C. Sierra, N. R. Jennings, P. Noriega, and S. Parsons, “A framework for argumentation-based negotiations,” in M. P. Singh, A. Rao, and M. J. Wooldridge, (eds.), *Intelligent Agents IV*, Springer: Berlin, Germany, pp. 177–192, 1998.
 52. M. P. Singh, “A social semantics for agent communications languages,” in F. Dignum, B. Chaib-draa and H. Weigand, (eds.), *Proceedings of the IJCAI-99 Workshop on Agent Communication Languages*, Berlin, Germany: Springer, 2000.
 53. K. Sycara, “Persuasive argumentation in negotiation,” *Theory and Decision*, vol. 28, pp. 203–242, 1990.
 54. G. L. Urban, J. R. Hauser, and J. H. Roberts, “Pre-launch forecasting of new automobiles,” *Management Science*, vol. 36, no. 4, pp. 401–421, 1990.
 55. D. N. Walton and E. C. W. Krabbe, *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*, SUNY Series in Logic and Language, State University of New York Press: Albany, NY, USA, 1995.
 56. Y. Wind and G. L. Lilien, “Marketing strategy models,” in J. Eliashberg and G. L. Lilien, (eds.), *Handbooks in Operations Research and Management Science: Volume 5: Marketing*, North-Holland, Amsterdam, The Netherlands, pp. 773–826, 1993.
 57. D. R. Wittink and P. Cattin, “Commercial use of conjoint analysis: An update,” *Journal of Marketing*, vol. 53, pp. 91–96, 1989.
 58. M. Wooldridge, “Semantic issues in the verification of agent communication languages,” *Journal of Autonomous Agents and Multi-Agent Systems*, vol. 3, no. 1, pp. 9–31, 2000.
 59. M. Wooldridge and S. Parsons, “Languages for negotiation,” in W. Horn, (ed.), *Proceedings of the Fourteenth European Conference on Artificial Intelligence (ECAI 2000)*, Berlin, Germany: IOS Press, pp. 393–397, 2000.