R&D MANAGEMENT

Learning patterns in early stage R&D projects: empirical evidence from the fibre raw material technology project in the Netherlands

Maryse M. H. Chappin[1,*], Jan Faber[1] and Marius T. H. Meeus[2,3]

[1]Copernicus Institute of Sustainable Development, Innovation Studies Group, Utrecht University, P.O. box 80115, 3508 TC, Utrecht, The Netherlands. m.m.h.chappin@uu.nl; J.Faber1@uu.nl
[2]Department of Organisation Studies, Tilburg University, P.O. Box 90153, 5000 LE, Tilburg, The Netherlands
[3]Center for Innovation Research, Tilburg University, P.O. Box 90153, 5000 LE, Tilburg, The Netherlands. m.t.h.meeus@tilburguniversity.edu

Past research has reported that learning processes in early stage R&D are either chaotic, or absent. We challenge this finding by elaborating Van de Ven et al.'s trial-and-error learning model and explore an alternative conceptualization. We explored the combinations of positive and negative outcomes and action course continuation and modification. We use data gathered in an R&D setting of a 4-years pre-competitive knowledge generation project in the Dutch paper and board industry. Whereas the Van de Ven and Polley (1992) approach applied on our data also would lead us to conclude that 'no learning' would happen, our decomposed model identified three distinct learning patterns: (1) a virtuous pattern of positive outcomes resulting in continuations of action courses; (2) a vacuous pattern of negative outcomes resulting in modifications of action courses; and (3) a verification pattern of positive outcomes resulting in modifications of action courses. We observed the virtuous and verification patterns during the first 2 years and virtuous and vacuous learning in the second 2 years. These results might be useful for R&D managers since they provide insight into how an early stage R&D project can develop and where managers might intervene and adjust action courses.

1. Introduction

Outcomes of innovation efforts are in general uncertain (Ingvarsson Munthe et al., 2014). Especially in R&D trajectories, it is most of the times upfront hard to predict whether their subprojects will be successful (Rice et al., 2001; Paul et al., 2010). Innovation is not a linear process, but a process of trial-and-error (Paul et al., 2010) and experimentation (Bingham and Davis, 2012). Van de Ven and colleagues (1999) conceptualize R&D and innovation as a trial-and-error learning process consisting of iterative sequences of actions and outcomes. This conceptualization builds on the research of Cyert and March (1963) and was the core theoretical model guiding the Minnesota Innovation Research Program (MIRP) (Van de Ven et al., 1999). In their empirical research Van de Ven et al. (1999) discern three stages in innovation trajectories: an initiation stage (early stage), an expansion stage and a contraction stage (late stage). The trial-and-error learning models of Van de Ven et al. (1999) turned out to be empirically valid only in the later stages, mainly the contraction stage, but not in the early stages of innovation trajectories. The inferences made by Van de Ven et al. (1999) and Cheng and Van de Ven (1996) about learning in early stage R&D as being 'chaotic', are the starting point of this paper. The front-end of innovation processes are featured by noisy action-outcome nexus seemingly determining their fuzziness. We ask the question: if the fuzzy front end is a chaotic process, surrounded by noise, and equivocality, what would that imply for the conceptualization and measurement of action-outcome links – the cornerstone of trial-and-error learning models?

In their studies, Van de Ven and co-authors (e.g. 1999) conceptualize and measure successive trial-and-error learning sequences as composite variables of differences in positive and negative outcomes and in continued and modified actions over time. The assumption behind these measures is that 'rational' researchers count positive and negative outcomes and calculate a 'net' outcome of past actions by subtracting negative from positive actions. The net outcome informs what action course to continue and what action course to change and result in a net action course for the next period of time. But these composite measures 'remove' the direct link between outcome and action. Moreover, the measurement aggregates learning processes, and therefore, conflates learning from success, and learning from failure. Gong et al. (2017) show that it is useful to differentiate these types since they have distinct effects on learning. So, the conceptualization and measurement of Ven de Ven and colleagues dismisses the possibility that learning at the project level consists of a series of simultaneous, but separate learning cycles.

These arguments require an alternative conceptualization and measurement to identify underlying learning patterns in early stage R&D. In early stage R&D projects, just after the start of R&D projects, much information is uncertain or missing (Stevens, 2014; Gouvêa De Oliveira et al., 2015) and outcomes will be very volatile. Early stage R&D resembles a search process in which researchers assess the quality and relevance of outcomes of various R&D activities (Courtney and Lovallo, 2004), such as experiments and field studies. Research outcomes are judged in multiple ways: usefulness, robustness and their guidance to plan further research actions. Hence, most outcomes have potential merits given certain goals in the R&D process. Even though the fuzzy front end of innovation has been studied extensively, insights into how actions as responses to R&D outcomes are differentiated over time, and to what extent these outcomes do indeed provide reliable and robust guidance for subsequent action patterns in early stage R&D is missing. Several scholars focused on the management of the fuzzy front end (Kim and Wilemon, 2002; Colombo et al., 2015) and identified success factors, such as leadership (Kach et al., 2012; Robbins and Gorman, 2015), communication (Verworn, 2009) and learning strategies (Stevens, 2014), which managers can use to reduce fuzziness. However, what specific learning patterns can be expected in early stage R&D is less clear. We pursue a more fine-grained approach for the identification of learning processes by decomposing the composite trial-and-error learning variables of Van de Ven and co-authors (1999) and exploring a larger range of action – outcome patterns building on the original concepts and measures applied. Our main research question is: *what learning patterns can be identified in early stage R&D and do they change over time?*

This paper is structured as follows. Next we explain our conceptualization of learning patterns in early stage R&D. Then we describe our data and methods applied to assess learning patterns, followed by the results. The paper ends with a discussion and conclusion.

2. Elaborating the trial-and-error learning model

2.1. *The evolution of learning theory in the behavioural theory of the firm*

Our theoretical model builds mainly on learning theory derived from the behavioural theory of the firm (Cyert and March, 1963). Over time, the two broad labels to categorize theories on learning in organizations, i.e. the behavioural and the cognitive approach, have been merged in Levitt and March's (1988) cognitive learning theory. The behavioural approach stems from Cyert and March (1963) and revolved around a stimulus–response behavioural approach to learning, in which external shocks or new requirements act as stimuli that require change in routines, protocols, that are stored in standard operating procedures serving as organizational memory. Organizations are able to learn from combinations of events, and adapt actions by its outcomes. Any action that leads to a preferred state at one point is likely to be used in the future again. Learning is a form of reactive adaptation in line with stimulus–response – operant conditioning – learning principles, without cognitive or knowledge related learning. March and Olson (1976) elaborate the concept of organizational learning to integrate cognitive structures. Especially the inclusion of beliefs, frameworks, paradigms, codes and knowledge develops the initial stimulus – response framework into a stimulus–interpretation–response framework as known in cognitive learning theory. March (1991) extends the individual learning model to a social learning model, interestingly without referencing social cognitive learning theory of Bandura. He opts for a mutual learning model (March, 1991) in which workers are socialized with routines and beliefs by observing peers, and colleagues within departments, projects and business units. In social cognitive learning theory (Bandura, 2001) self-efficacy motivates agency, because ultimately humans can produce desired results while forestalling detrimental outcomes. In social cognitive theory this efficacy belief makes performance feedback on past actions the main driver of future behaviour.

2.2. *The trial-and-error learning model*

Trial-and-error learning consists of an iterative sequence of actions and outcomes (Cyert and March, 1963; Levitt and March, 1988) that reinforce or inhibit prior actions, because of feedbacks about its effectiveness derived from either positive or negative outcomes. In order to learn, researchers initiate an action course A to achieve positive outcomes. If the expected positive outcomes occur, action course A will be continued until the ultimate goals are achieved. In case of negative outcomes inducements emerge to change the course of action, and to search for an alternated course B that could result in positive outcomes. Action courses will persist as long as expected positive outcomes of actions are found; an action course will be modified when negative outcomes are experienced (Van de Ven and Polley, 1992; Van de Ven et al., 1999). So, as also indicated by Gong et al. (2017) success and failure (different types of outcomes) have distinct learning effects. Innovation processes consist of series of these internal loops linking actions to distinct goals, tasks and outcomes. The mechanism describing this pairing of outcomes and actions has been labelled the *Law of Effect* (Thorndike, 1911).

The MIRP project headed by Andrew Van de Ven, explored trial-and-error learning processes in a variety of R&D settings (Zahra, 2016). For the initiation stage, Van de Ven and Polley (1992) report that no learning occurred. They provide a variety of arguments for this finding: overrating purposefully the outcomes of the project, high team turnover which hampered the accumulation of past experience and organizational memory, proliferation of action courses loosely connected to the core project idea, and lack of recognition of setbacks (negative outcomes). These findings and explanations are indicative of an absence of decision rationality in early stage R&D, and that no reinforcement from past actions outcomes could make the organizations learn. Cheng and van de Ven (1996) provide additional reasons for the absence of trial-and-error learning in the earlier stages. First, they argue that both actions and outcomes turned out to be chaotic in the beginning because of unclear preferences and ambiguous goals and only develop into a periodic pattern in the ending period. Second, they argue that because learning patterns evolve over time, as the beliefs researchers started out with only gradually develop into reproducible and objectified knowledge (Cheng and Van de Ven, 1996). In sum, the research of Van de Ven, Polley and Cheng make a case for the low explanatory value of the trial-and-error learning model in early stages of innovation. However, they take their measures and model for granted, and do not reflect on alternative conceptualizations and model to assess the occurrence of trial-and-error learning. This lacuna is the niche that justifies our study.

2.3. *Towards an alternative trial-and-error model*

The findings of Van de Ven and co-authors (e.g. Van de Ven and Polley, 1992; Cheng and Van de Ven, 1996; Van de Ven et al., 1999) imply that the core explanatory mechanism in trial-and-error learning models have to be played down in the fuzzy front end of innovation. Thorndike's (1911) Law of Effect states that people choose behaviours contingent on their outcomes. However, there is a strong assumption in the Law of Effect on the clarity as well the assessment of goals, and hence the directionality of learning behaviours. One can raise the issue, whether learning in *early stage R&D* is structured, as assumed in the Law of Effect.

Cheng and Van de Ven (1996) report that at the start R&D projects often lack clear goals, and conditions. Much information at this stage is either uncertain or absent (Gouvêa De Oliveira et al., 2015). As a consequence the choices and decisions researchers have to make are driven by the 'highest level of fuzziness' (Stevens, 2014). Instead of having specific projects goals, people explore possibilities and opportunities within R&D projects, often in concurrent work on several parts of those projects. As to the nature and targets of learning in early stage R&D many planned actions are grounded in beliefs (Courtney and Lovallo, 2004) and prior practical experiences of the partners involved (Murray, 2004). This results in an inherently ambiguous setting, in which many actions might not at all deliver the expected outcomes, or even no discrete outcomes, as the results of some actions can be indeterminate (Van de Ven and Polley, 1992). Therefore, early stage R&D, need multiple iterations and experimentation to replicate outcomes, in order to disambiguate both the setting and goals. During these iterations, researchers start to develop an understanding of what it means when re-testing certain findings for robustness or sensitivity to certain contingencies. In this setting of high ambiguity positive and negative feedback might not work in the same way as it does with learning from routinized action-outcome linkages in production tasks, especially because intended goals are revisited, revised, updated both up- and downward.

Researchers work concurrently on several constituting parts of early stage R&D projects and, when time progresses they will produce outcomes. These outcomes can induce either continuation, or modification of prior action courses, but such inferences are not clear up front, which results in ambiguity in learning patterns. This has important implications for the conceptualization and the measurement of action – outcome sequences. The composite measures of action courses and outcomes of Van de Ven et al. (1992) do not allow for the identification of learning processes, because it masks the differences between action course continuations after positive outcomes as well as action course modifications after negative outcomes. So, we infer a deficit in this measure of reinforcement learning pattern as it hides which outcome-action course combination prevails, while the coincidence of multiple action-outcomes combinations cannot be discerned from the result obtained. To solve this conceptual and measurement issue, we propose, an alternative conceptualization and measurement of learning patterns that allows us to identify distinct learning patterns that occur simultaneously. Opposite to Van de Ven et al. (e.g. 1999), we do not aggregate the outcomes, but separate both positive and negative outcomes, as well as the two types of action courses. Accordingly, every possible combination of positive and negative outcomes and the subsequent action course of continuation and modification being either positively or negatively related is taken into account and is conceptualized as a distinct learning pattern. This means that we rigorously reconceptualize learning as: the alternation and accumulation of distinct learning patterns over time. In doing so, we play down Thorndike's Law of Effect and leave more space for belief-based learning and learning from failure.

In order to gain insight into the relevance of our conceptualization, we will compare the results when using composite measures of Van de Ven et al., next to decomposed count measures of positive and negative outcomes and decomposed count measures of subsequent continued and modified action courses. It is, however, impossible to indicate beforehand which specific learning patters will be identified. They will be further explored in the empirical research carried out.

3. Data and methods

3.1. Research context

In this paper, we focus on early stage R&D in the Dutch paper and board industry. The production of paper is energy intensive. Moreover, in the Dutch context waste and also waste water are important topics for the industry (Chappin, 2008). Because of these environmental challenges, process innovation has been prominent in this industry over time (Chappin et al., 2007). Knowledge development supporting process innovation has been supported by the Dutch Centre of Competence Paper and Board (KCPK) since 1998. In order to explore learning patterns, we investigated one of the early stage R&D projects of the KCPK. The project we studied is called *Fibre raw material technology for sustainable production of paper and board*. It is a pre-competitive knowledge generation project, funded by the Dutch Ecology, Economy and Technology funding scheme. The environment and the reduction of environmental impact were important drivers for the start of this research project. The project lasted 4 years (early 2000 until the end of 2003) and was a collaboration of 13 organizations: 2 paper and board producers, 5 paper industry suppliers, 3 universities, 2 research institutes and the KCPK.

Four sub-projects were executed by working groups A–D. Each sub-project team had its own specific objective, but there was exchange between the teams. Teams also needed to collaborate and share knowledge, so they had a coordinator and regular meetings, a steering group and project management teams. The project coordinators governed the integration of tasks, planning and subprojects.

The general aim was 'to develop new innovative technologies for the processing of fibre raw materials for paper and board production' (Van Kessel and Westenbroek, 2004, p. 9).

The following five objectives were set (Van Kessel and Westenbroek, 2004):

- Closure of the fibre cycle: less waste, less water pollution and higher energy efficiency, mostly by re-use of waste and water.
- More effective use of fibre raw materials
- Improved control and efficiency of fibre processing
- Improved product quality
- Strengthening and broadening the knowledge infrastructure of the Dutch paper and board industry

In order to achieve these objectives working groups set out to develop fundamental knowledge on changes in the features of fibre during paper production, and on technologies to conserve or upgrade the fibres.

3.2. Data collection

The first step in the data collection process was to build a database that consisted of a chronological listing of all incidents that took place within the project. To build the database we used the following sources:

- Minutes of meetings of working groups, project management team and steering group
- Reports written for the project
- Presentations within the project
- Archive KCPK
- Face-to-face communication for clarification

The minutes describe the judgement and choices of the researchers, practitioners and managers in their own language. Using these minutes, biases induced by the judgement of external observers or the notes of writers who lack understanding of the research issues at hand are avoided. The minutes are written by the R&D people involved for their own use to track the progress of the project. This written documentation was made real time during the project. This is an important asset, as we do not have to rely on the memories of the project members about past experiences, which are possibly subject to memory-decay. Nevertheless, individual biases of R&D people might affect their judgement and evaluation of both action courses and outcomes. We tried to curtail this limitation by reading additional sources and crosschecking the minutes with research reports, research presentations, archival material and oral communications.

3.3. Data preparation

The data preparation for multivariate analyses consisted of coding the chronological list of incidents that referred to actions and/or outcomes pertinent to the knowledge content of the project. Table 1 shows our operationalization.

In order to check for subjectivity of the coding, three other coders coded a different subset of the database. There was a large agreement between the codings of the main coder and the other coders (Cohen's Kappa > 0.8). Differences in coding were discussed and when necessary adjusted.

To compare our findings with Van de Ven and Polley (1992) we also applied a similar operationalization: outcomes are calculated as a composite measure of the number of positive outcomes minus the number of negative outcomes per 2 months period[1]; action continuation–modification is also calculated as a composite of the number of continuing actions minus the number of modified actions per 2 months period. In the second step of the analysis, we use the four decomposed variables: positive and negative outcomes, action continuations and action modifications per 2 months period.

As the outcomes and the subsequent actions are reported in the same sections of the minutes, and thus at the same moment in time, we do not have time dependent effects between the dependent and independent variables in our analyses. Nevertheless, a causal structure between outcome and follow-up action can be discerned as the minutes report upon the judgement of the nature of the outcomes and which actions are needed next.

3.4. Data analysis

The aforementioned two monthly observations result in categorical variables with only a few categories of count data. The correlations between these variables have been estimated by means of the polychoric correlation coefficient (Olsson, 1979), which represent Pearson correlations of the standardized, normally distributed variables underlying the categorical variables representing their discrete realizations. Accordingly, these correlations are used as input for multivariate analyses. Since, we are interested in identifying periodic patterns we can use linear models (Cheng and Van de Ven, 1996).

R&D projects have a fixed time limit. However, how learning processes unfold over time is important

Table 1. Constructs measurement

Construct	Dimensions	Indicators and measurement	Example Incidents*
Action Course related to the research project	Continuation	If a documented incident displays a continuation of an course such as an experiment, literature study, desk top research → This incident is coded as ActionCourse_Conti = 1	Partner x will continue experiments to look at y
	Modification	If a documented incident displays a modification in the action course such as an experiment, literature study, desk top research. → This incident is coded as ActionCourse_Modi = 1	Before a detailed work plan will be made member x first wants to make an inventory about y
Outcomes related to the research project	Positive	If a documented incident displays a result of an experiment, literature study or desk top research that is positive or in line with the expectations → This incident is coded as Outcome_Pos = 1	It was concluded that for process x, y is an interesting additional process variable
	Negative	If a documented incident displays a result of an experiment, literature study or desk top research that is negative or not in line with expectations → This incident is coded as Outcome_Neg = 1	Partner x concluded that y is not an option to study
	Mixed	If a documented incident displays a result of an experiment, literature study or desk top research that is positive as well as negative → This incident is coded as Outcome_Pos = 0.5 and Outcome_Neg = 0.5	x led to a better correlation than y however y still gives significant higher values in the case of z

*Due to confidentiality, specific content of the incidents has been made anonymous.

but generally left unexplored (Weber and Berthoin Antal, 2003). Deadlines make individuals and teams more aware of time urgency, which is typically experienced when the mid-point comes closer (Waller et al., 2002). This is called mid-point transition (Gersick, 1989). Figure 1 provides an overview of the actions and outcomes over time. At period 12 (grey area Figure 1), we can observe a midpoint transition (Gersick, 1989), showing an increase in work activity at the midpoint of the project. Consequently, we decided to divide the project in two periods with equal numbers of observation.

For each period, we conducted regression analyses. In the first analysis, we use the composite variables similar to Van den Ven and Polley (1992). In the second analysis, we used our decomposed measures in the analyses. As the dependent variables action course continuation and modification are not

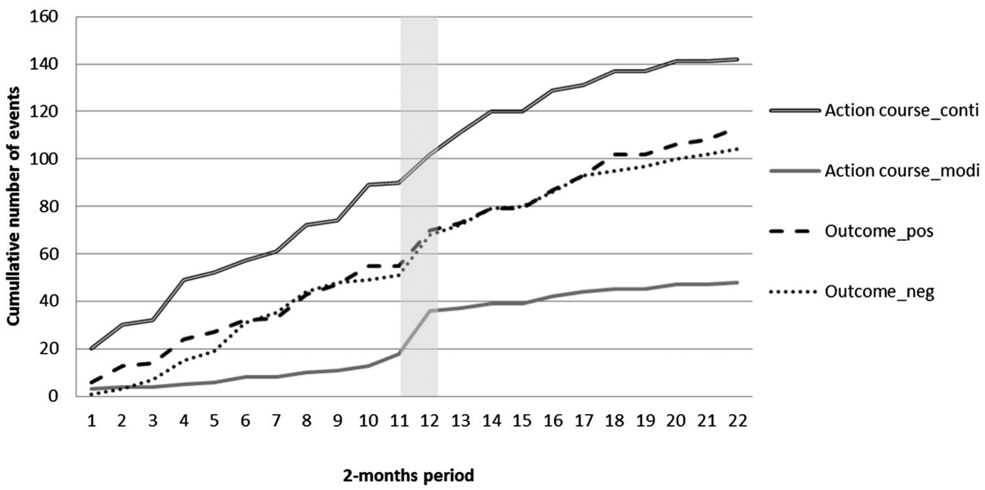

Figure 1. Cumulative number of actions and outcomes.

unrelated and both dependent on the same independent variables, their regression equations are simultaneously estimated in one linear model by means of the maximum likelihood method in LISREL 8 (Jöreskog and Sörbom, 1996). As we use panel data, we control for fixed working group effects (by including the working groups as dummy variables) and autocorrelation in the dependent variables (by including the one period lag of the dependent variable).

4. Results

4.1. Descriptive statistics

Table 2 presents descriptive statistics. Over the complete duration of the project we observed 142 continuations and 48 modifications of action courses. The total number of references to positive and negative outcomes are 113 and 104. Comparing the two periods shows that almost two-third of the action course continuation took place in the first period. For modification of the action course this is the other way around. Regarding the outcomes (positive and negative), we see that their references are more or less evenly distributed over both periods.

4.2. Results regression analyses for the composite measures

Table 3 shows the results of the regression analyses based on the composite measures. The positive coefficients for the effect of outcome on action course provide an indication of trial-and-error learning. However, none of these relations are statistically significant. Accordingly, there seems to be no evidence of the systematic presence of trial-and-error learning in both periods of this R&D project.

4.3. Results regression analyses for the decomposed measures

Table 4 displays the results of the analysis with decomposed variables.[2]

Models 1 and 2 show the results for the first period (years 1–2) and the models 3 and 4 for the second period (years 3–4). All models are statistically

Table 2. Descriptive statistics

	Period I						Period II					
	n	Min	Max	Sum	Mean	Std. Dev.	***n***	Min	Max	Sum	Mean	Std. Dev.
ActionCourse_Conti	44	0	11	90	2.05	2.72	44	0	7	52	1.18	2.18
ActionCourse _Modi	44	0	4	18	0.41	0.79	44	0	8	30	0.68	1.61
ActionCourse	44	−4	10	[1]	[1]	[1]	44	−3	7	[1]	[1]	[1]
Outcome_Pos	44	0	6	55	1.25	1.53	44	0	8	58	1.32	2.02
Outcome_neg	44	0	5	51	1.16	1.48	44	0	7	53	1.21	1.72
Outcome	44	−4	5	[1]	[1]	[1]	44	−3	6	[1]	[1]	[1]

This descriptive does not make sense for this variable since it is a composite measure.

Table 3. Estimated standardized regression coefficients based on composite measures

	Years 1 and 2	Years 3 and 4
Independent variables	Action Course (t)	Action Course (t)
Outcome(t)	0.376	0.031
ActionCourse(t−1)	−0.349	−0.290
WorkingGroup_B+	−0.494	0.270
WorkingGroup_C+	−0.073	0.059
WorkingGroup_D+	−0.396	−0.382
R^2	0.326	0.307
adj. R^2	0.237	0.216
F-value	3.676***	3.367**

+WorkingGroup_A is the reference category.
*$P \leq 0.10$.
**$P \leq 0.05$.
***$P \leq 0.01$ (two-tailed).

significant and explain a substantial part of the variance in Action Course Continuation (ACC) and Action Course Modification (ACM).

In the *first period*, we see that researchers only systematically react to positive outcomes. We find a positive effect on the continuation of actions in model 1 ($b = 1.040$; $P < 0.01$) as well as the modification of actions in model 2 ($b = 1.061$; $P < 0.01$). The coefficients for the effect of negative outcomes on action course continuation and action course modification are not statistically significant. Furthermore, independent of previous outcomes working groups B, C and D are significantly more inclined to modify their action courses than working group A due to fixed working group effects of differences in, for example, prior knowledge, prior practical experience, age-related risk attitudes, etc.

For the *second period,* we observe the same positive effect of positive outcomes on the continuation of actions in model 3 ($b = 1.578$; $P < 0.05$). However, the positive effect of positive outcomes on the modification of actions is no longer significant. In this period, we also observe a systematic reaction to negative outcomes. Model 4 reveals that when the number of negative outcomes increases it is more likely that the number of modified action courses also increases ($b = 0.844$; $P < 0.1$). We also see a positive coefficient for the relation between negative outcomes and action course continuation, but this effect is not statistically significant. Fixed working group effects are absent in this period.

Our findings are indicative of three different learning patterns illustrated with an example derived from the research project materials. These examples should provide a better understanding of the specific learning patterns.

The first learning pattern is based on the idea that one continues an action if the outcomes are perceived positive. We label this as *the virtuous cycle*. As recycling was the core topic in the knowledge generation project, researchers developed a laboratory recycling procedure based on knowledge obtained from scientific literature and on the experience of the participating paper and board factories. Researchers first conducted a few trials that showed a decrease in fibre potential after several recycling loops. In the minutes, a workgroup reports that these results of the trials '*confirms that the procedure used is effective*'. Based on this positive evaluation of the recycling procedure, they decided to continue working with it.

The second learning pattern that we identified was the modification of actions due to positive outcomes. We label this *the verification cycle*. One of the tasks of the knowledge generation project was to research the mechanisms of degradation of fibre polymers during recycling. The expected differences in the paper properties as a consequence of ageing were confirmed. The researchers decided that in addition to the 'normal' judgement of these findings, other measurements were required. It was announced in the minutes that they were planning

Table 4. Estimated standardized regression coefficients based on decomposed measures

	Years 1 and 2		Years 3 and 4	
	Model 1	Model 2	Model 3	Model 4
Independent variables	ACC(t)	ACM(t)	ACC(t)	ACM(t)
Outcome_Pos(t)	1.040***	1.061***	1.578**	0.598
Outcome_Neg(t)	−0.371	0.095	0.085	0.844*
ActionCourse_Conti(t−1)	−0.096	na	0.703	na
ActionCourse_Modi(t−1)	na	0.819	na	0.358
WorkingGroup_B+	0.384	1.683*	−0.653	−0.293
WorkingGroup_C+	0.593	0.874*	−0.877	−0.509
WorkingGroup_D+	0.086	0.733**	0.086	0.195
R^2	0.692	0.392	0.722	0.945
Adj. R^2	0.642	0.293	0.677	0.936
F-value	13.855***	3.976***	16.016***	105.955***
Chi-square/df	5.016/3		0.467/3	

na: means not applicable.
ACC(t) = Action Course Continuation(t); ACM(t) = Action Course Modification(t).
+Working Group A is the reference category.
*$P \leq 0.10$.
**$P \leq 0.05$.
***$P \leq 0.01$ (two-tailed).

for additional measurements in order to have a better judgement.

The final pattern we identified was that prior negative outcomes feed forward into more modification of action courses. We labelled this the *vacuous learning pattern*. We observed this pattern for instance in the context of washing trials. When paper is being recycled fines and ash levels increase, which influence the paper and pulp properties, a possible effect being a decrease in porosity of the paper. Washing is a possible solution. The researchers conducted several pilot washing trials that showed that the pulp properties changed quite a bit. The findings made clear that more information was needed, and hence they performed additional and different washing trials.

A comparison of learning patterns between the two periods shows that the virtuous cycle occurs in both periods. Next, we also observe a switch in dominant learning patterns. The verification pattern withers away in the second period, while observing the emergence of vacuous cycle in this period. We interpret these findings as a result of project progression and accumulation of experience of researchers with their projects. Their experience will make them more confident to modify their behaviour as a reaction to negative outcomes.

In sum, the analyses based on the composite measures did not provide systematic evidence of trial-and-error learning in both periods of this early stage R&D project. The analyses based on the decomposed measures, however, identified three different learning patterns: a virtuous learning pattern and a verification learning pattern in the first period and a virtuous learning pattern and a vacuous learning pattern in the second period.

5. Discussion and conclusion

5.1. Theoretical implications

We took the non-findings on trial-and-error learning in early stage R&D, i.e. in the fuzzy-front-end of innovation as our starting point. Learning in early stage R&D has been characterized as chaotic. Our aim was to increase the understanding of the core mechanism – the association of two types of action outcomes, and subsequent action courses – underlying the trial-and-error learning in early stage R&D. We used the following research question: *what learning patterns can be identified in early stage R&D and do they change over time?*

We elaborated the trial-and-error learning model into more fine-grained approach. Instead of using composite variables for action courses and outcomes, we decomposed the action courses into separate variables for modification and continuation of actions and for negative and positive outcomes. Consequently, we explored which learning patterns could be identified.

Our findings replicate the original trial-and-error learning model tests (e.g. Van de Ven et al., 1999). However, the models using our decomposed variables revealed three possible learning patterns and our findings suggest a changing nature of the learning patterns over time as well.

First, we observed in both periods continuation of the action course in case of positive outcomes, which we labelled *virtuous learning pattern.* This virtuous learning is the core process of the trial-and-error model (Van de Ven et al., 1999). The virtuous pattern implies that positive outcomes are self-reinforcing, and generate convergence in knowledge outcomes and reinforce prior beliefs in actions taken. In an early stage R&D project, engineers and scientists build on these virtuous patterns as they carry their project to the next stage and ensure future investments in the R&D. The modification of the action course, while observing positive outcomes, is the second learning pattern that we observed in the first period and we labelled this as the *verification learning pattern*. Such behaviour is not adaptive in the way suggested by the Law of Effect, as satisficing outcomes are verified, re-tested, under different conditions, either to optimize outcomes or to rule out alternative positive outcomes. This logic is often applied in prototyping settings, to quickly learn about alternative search routes that seek to reveal technological opportunities, outside the mainstream (Brown and Eisenhardt, 1997). Finally, we observed in the second period a third pattern of negative outcomes feeding forward to the modification of action courses, which we labelled the *vacuous learning pattern*. This pattern is also in line with the original trial-and-error learning as developed in earlier research on trial-and-error learning (e.g. Van de Ven and Polley, 1992). Other learning patterns that can be identified with the decomposed model are the following. If positive outcomes have a negative effect on continuation as well as modification of action courses, this is indicative of results of early stage R&D becoming satisfactory in the light of goals set. If negative outcomes have a negative effect on continuation as well as modification of action courses, this is indicative results of early stage R&D providing no further clues of how to proceed and the project has failed. If negative outcomes have a positive effect on the continuation of action courses, this is indicative of behavioural persistence in accordance with the initial plan, or with belief learning.

With respect to the changes in learning patterns over time we draw several inferences. First, it turns out that the pattern of vacuous learning is insignificant in the first period as the guidance for follow-up actions derived from negative outcomes might be rather limited in the beginning of a project. In that stage researchers rely mainly on their beliefs, and prior experience, because robust, validated knowledge and experience built up within the project is still absent. Accordingly, researchers are likely to respond mainly to outcomes that confirm these beliefs. In the second period, researchers gradually develop more robust and reliable insights that feeds into more confidence in action course selection. This enables them to adjust their behaviour based on negative findings. This might explain the insignificance of the vacuous cycle in the first period.

Whereas Cheng and Van de Ven (1996) and Dooley and Van de Ven (1999) concluded that early stage R&D tends to be 'chaotic', our findings show that early stage R&D entails more than just reinforcement learning as suggested by the earlier conceptualization of trial-and-error learning of Cyert and March (1963). Our findings show that in this early stage R&D project, researchers found different ways out of the 'chaos' that occurs in early stage R&D via simultaneously unfolding virtuous, verification and vacuous learning patterns. This finding opens up new venues for identifying learning patterns, and dynamics in these learning patterns albeit early stage R&D.

5.2. Limitations and recommendations for future research

Our study has limitations. The empirical results obtained in this study are limited to this setting. Due to the labour intensive process of data collection and analysis, we studied only one, but successful and unique early stage R&D project. In the learning literature, similar studies that explored cases in such detail remain rare. The MIRP project of Van de Ven and colleagues is one of the exceptions to the rule (Zahra, 2016). We cannot and will not claim that the observed learning patterns occur in other projects as well. However, the fact that the findings of Van de Ven et al. are replicated with our data is reassuring, as it extends the external validity of their findings as much as it extends the external validity of our findings. Furthermore, this replication also adds to the internal validity of our decomposed model, which is a more elaborated version of Van de Ven et al.'s trial-and error learning model based on an analytical and theoretical decomposition of which trial-and error learning is a part. This allows us to conclude that our findings are not simply an artefact of our own new measurement model only but also bear empirical and theoretical relevance. Nevertheless, as always, a recommendation for further research is to test the model of singular outcomes and action courses in a large number of early stage R&D projects. Furthermore, studies of learning patterns in completely successful and failed projects might be relevant, as we only examined a rather but not completely successful research project. Similarly, it would be interesting to explore whether the success or failure of later stage R&D projects can be explained by our more fine-grained approach.

5.3. Managerial implications

Our reconceptualization and decomposition of the trial-and-error learning model offers a straightforward set of concepts and measures to help R&D and project managers to identify learning patterns early on in R&D projects. Our findings highlight the importance of keeping track of the actions project teams take as a response to negative or positive outcomes, and offer insight in the timing of these actions. This inference is evidenced by the fact that learning patterns changed between the first and the second time period. These temporal changes suggest that project and R&D managers might consider to start by focusing on continuation when positive outcomes are observed (virtuous patterns) and to alternate reactions to negative outcomes (vacuous patterns) over time. At the start negative outcomes cannot offer much guidance on action courses, whereas later on in a project such outcomes definitely reduce the set of potential actions.

All in all, every learning pattern seems to matter for project management and to have a function in motivating R&D project members, preventing time and budget overruns, while securing knowledge gains, and closing down too divergent research options. In this way, our exploratory research offers a substitute to replace Van de Ven et al's 'chaos' in early stage R&D, and the fuzzy front end of innovation.

Acknowledgements

We would like to thank the Kenniscentrum Papier en Karton, and in particular Arie Hooimeijer and Annita Westenbroek, for their openness and willingness to collaborate on this project. Furthermore, we would like to thank Andrew van de Ven, Kay Schoonhoven, Amalya Oliver, Fariborz Damanpour and Joris Knoben for their valuable comments on earlier versions of this paper. Finally, we would like to thank the participants of EGOS SWG02 2008 for their comments.

References

Bandura, A. (2001) Social cognitive theory: an agentic perspective. *Annual Review of Psychology*, **52**, 1–26.

Bingham, C.B. and Davis, J.P. (2012) Learning sequences: their existence, effect, and evolution (En 10.5465/amj.2009.0331). *Academy of Management Journal*, **55**, 611–641.

Brown, S.L. and Eisenhardt, K.M. (1997) The art of continuous change: linking complexity theory and time-paced evolution in relentlessly shifting organizations. *Administrative Science Quarterly*, **42**, 1.

Chappin, M.M.H. (2008) *Opening the Black Box of Environmental Innovation. Governmental Policy and Learning in the Dutch Paper and Board Industry*. Utrecht: Utrecht University.

Chappin, M.M.H., Meeus, M.T.H., Hekkert, M.P., and Vermeulen, W.J.V. (2007) Dynamic perspective on the relation between environmental policy and eco-efficiency: the case of wastewater treatment, waste and energy efficiency in the Dutch paper and board industry. *Progress in Industrial Ecology*, **4**, 19–40.

Cheng, Y.-T. and Van de Ven, A.H. (1996) Learning the innovation journey: order out of chaos? *Organization Science*, **7**, 593–614.

Colombo, G., Dell'Era, C., and Frattini, F. (2015) Exploring the contribution of innovation intermediaries to the new product development (NPD) process: a typology and an empirical study. *R&D Management*, **45**, 126–146.

Courtney, H. and Lovallo, D. (2004) Bringing rigor and reality to early-stage R&D decisions. *Research-technology Management*, **47**, 40–45.

Cyert, R.M. and March, J.G. (1963) *A Behavioral Theory of the Firm*. Englewood Cliffs, NJ: Prentice-Hall Inc.

Dooley, K.J. and Van De Ven, A.H. (1999) Explaining complex organizational dynamics. *Organization Science*, **10**, 358–372.

Gersick, C.J.G. (1989) Marking time: predictable transitions in task groups (En 10.2307/256363). *Academy of Management Journal*, **32**, 274–309.

Gong, Y., Zhang, Y., and Xia, J. (2017) Do firms learn more from small or big successes and failures? A test of the outcome-based feedback learning perspective. *Journal of Management*, 1–23. Available at https://journals.sagepub.com/doi/10.1177/0149206316687641#articleCitationDownloadContainer.

Gouvêa De Oliveira, M., Rozenfeld, H., Phaal, R., and Probert, D. (2015) Decision making at the front end of innovation: the hidden influence of knowledge and decision criteria. *R&D Management*, **45**, 161–180.

Ingvarsson Munthe, C., Uppvall, L., Engwall, M., and Dahlén, L. (2014) Dealing with the devil of deviations: managing uncertainty during product development execution. *R&D Management*, **44**, 203–2016.

Jöreskog, K.G. and Sörbom, D. (1996) *LISREL® 8: User's Reference Guide*. Chicago: Scientific Software International Inc.

Kach, A., Azadegan, A., and Dooley, K.J. (2012) Analyzing the successful development of a high novelty innovation project under a time pressured schedule. *R&D Management*, **42**, 377–400.

Kim, J. and Wilemon, D. (2002) Focusing the fuzzy front–end in new product development. *R&D Management*, **32**, 269–279.

Levitt, B. and March, J.G. (1988) Organizational learning. *Annual Review of Sociology*, **14**, 319–340.

March, J.G. (1991) Exploration and exploitation in organizational learning. *Organization Science*, **2**, 71–87.

March, J.G. and Olsen, J.P. (1976) *Organizational Choice under Ambiguity In Ambiguity and Choice in Organizations*. Bergen: Universitetforlaget. pp. 10–23.

Murray, F. (2004) The role of academic inventors in entrepreneurial firms: sharing the laboratory life. *Research Policy*, **33**, 643–659.

O'Brien, R. (2007) A caution regarding rules of thumb for variance inflation factors. *Quality & Quantity*, **41**, 673–680.

Olsson, U. (1979) Maximum likelihood estimation of the polychoric correlation coefficient. *Psychometrika*, **44**, 443–460.

Paul, S.M., Mytelka, D.S., Dunwiddie, C.T., Persinger, C.C., Munos, B.H., Lindborg, S.R., and Schacht, A.L. (2010) How to improve R&D productivity: the pharmaceutical industry's grand challenge. *Nature Reviews. Drug Discovery*, **9**, 203–214.

Rice, M., Kelley, D., Peters, L., and Colarelli O'Connor, G. (2001) Radical innovation: triggering initiation of opportunity recognition and evaluation. *R&D Management*, **31**, 409–420.

Robbins, P. and Gorman, C.O. (2015) Innovating the innovation process: an organisational experiment in global pharma pursuing radical innovation. *R&D Management*, **45**, 76–93.

Stevens, E. (2014) Fuzzy front-end learning strategies: exploration of a high-tech company. *Technovation*, **34**, 431–440.

Thorndike, E.L. (1911) *Animal Intelligence*. New York: Macmillan.

Van de Ven, A.H. and Polley, D. (1992) Learning while innovating. *Organization Science*, **3**, 92–116.

Van de Ven, A.H., Polley, D.E., Garud, R., and Venkataraman, S. (1999) *The Innovation Journey*. New York: Oxford University Press.

Van Kessel, L.P.M. and Westenbroek, A.P.H. (2004) *Fibre Raw Material Technology for sustainable paper and board production*. Arnhem: Kenniscentrum Papier en Karton.

Verworn, B. (2009) A structural equation model of the impact of the 'fuzzy front end' on the success of new product development. *Research Policy*, **38**, 1571–1581.

Waller, M.J., Zellmer-Bruhn, M.E., and Giambatista, R.C. (2002) Watching the clock: group pacing behavior under dynamic deadlines (En 10.2307/3069329). *Academy of Management Journal*, **45**, 1046–1055.

Weber, C. and Berthoin Antal, A. (2003) The role of time in organizational learning. In: *Handbook of Organizational Learning and Knowledge*. New York: Oxford University Press Inc. pp. 351–368.

Zahra, S.A. (2016) Portrait of a research pioneer: Andrew Van de Ven. *Strategic Entrepreneurship Journal*, **10**, 413–429.

Notes

1. In order to estimate the models, time series needed to be created. Similar to Van de Ven and Polley (1992), different time series of all variables involved have been created and compared: of the dates of incidents, aggregated over 1 month, aggregated over 2 months, and aggregated over 3 months. As suggested by Van de Ven and Polley (1992:100), we chose the interval that provided us with the most substantively meaningful interpretation of the time series graphs and correlations. In our case that was the time series of incidents aggregated over 2 months.
2. Some of the larger unbiased estimates of regression coefficients, which were statistically insignificant, may have suffered from inflated estimates of the variances of their standard errors due to some large correlations (>0.70) among the independent variables specified, i.e. multicollinearity. However, leaving out some of the controls for fixed group effects and/or autocorrelation in the dependent variables, both necessary in the analysis of longitudinal panel data, or even one of the outcome variables (+/-) would lead to seriously biased estimates of the effects of the regression coefficients of the independent variables left in the analysis. So, the choice is between obtaining unbiased and possibly insignificant estimated regression coefficients or biased and possibly significant estimated regression coefficients. O'Brien (2007) convincingly argues to choose for the first option as we also did.

Maryse M. H. Chappin is an assistant professor of Innovation Studies in the Copernicus Institute of Sustainable Development at Utrecht University from where she received her PhD. Her work focuses on eco-innovation, networks and institutions, and on crowdsourcing for innovation.

Jan Faber is an associate professor of Innovation Studies in the Copernicus Institute of Sustainable Development at Utrecht University. His research and teaching activities are focused on technology related venturing. He currently supervises MSc and PhD students on topics related to technology based product and process innovations by established and new firms. Jan Faber has a background in Economic Geography (MSc, 1980) and Political Geography (PhD, 1987).

Marius T. H. Meeus is full professor of organization science at Tilburg University, from where he received his PhD. He is director of the Center for Innovation Research. His research interests include networks, learning, innovation and organizational adaptation.

