

**SYNCHRONIZING EXPLORATION AND EXPLOITATION:  
KNOWLEDGE CREATION CHALLENGES IN INNOVATION**

A Dissertation  
Presented to  
The Academic Faculty

by

Jennifer Bailey

In Partial Fulfillment  
of the Requirements for the Degree  
Doctor of Philosophy in the  
Scheller College of Business

Georgia Institute of Technology  
December 2013

**COPYRIGHT 2013 BY JENNIFER BAILEY**

**SYNCHRONIZING EXPLORATION AND EXPLOITATION:  
KNOWLEDGE CREATION CHALLENGES IN INNOVATION**

Approved by:

Dr. Cheryl Gaimon, Advisor  
Scheller College of Business  
*Georgia Institute of Technology*

Dr. Vinod Singhal  
Scheller College of Business  
*Georgia Institute of Technology*

Dr. Manpreet Hora  
Scheller College of Business  
*Georgia Institute of Technology*

Dr. David Ku  
School of Mechanical Engineering  
*Georgia Institute of Technology*

Dr. Stylianos Kavadias  
School of Business  
*Cambridge University*

Date Approved: November 15, 2013

*"There is the risk you cannot afford to take,  
and there is the risk you cannot afford not to take."*

*- Peter Drucker*

*For my parents*

*Rev'd. C. Evans and Dr. Barbara Bailey*

*and*

*for Betty and Grace-Ann.*

## **ACKNOWLEDGEMENTS**

The PhD, I have been told, is like an extreme roller coaster ride: You get to the top, look down, and then wonder, why you ever decided to do it. And then you get to the bottom and you are glad that you did. By no means has this journey been a smooth one, but it is definitely a journey I am glad that I embarked upon. And like all successful sojourns, it was made all the better with good guides and good companions. I owe a debt of gratitude to those without whom this would not been possible.

I am grateful to have had a superb guide in my advisor Dr. Gaimon. Throughout my PhD program, Dr. Gaimon has always given me her support, taking her time to patiently explain, correct and mentor. Through her academic and leadership achievements, she has provided a clear path, to demonstrate how others can also have a successful journey. And as I developed my research, she has consistently encouraged me to boldly explore my own new path.

I am thankful that Dr. Hora joined the faculty just as I entered the program. Immediately, he provided support to explore my interests, through an independent study. With him I have been able to examine the empirical perspective of my research. I thoroughly enjoyed our lively and thought-provoking discussions, but am most appreciative of his unwavering encouragement and support. Dr. Toktay's quiet but consistent mentoring came at all the critical times, and her vote of confidence was always there, even if my wavered. I am so happy that I had the opportunity to take Dr. Kavadias' NPD seminar and to have experienced first-hand his passion for his work, which is so contagious. His passion equally matched by Dr. Rothaermel, whose seminar fueled my interest in

ambidexterity. Together, with supply chain and sustainability seminars from Drs. Subramanian and Atasu I felt prepared with tools to begin the academic journey. Dr. Atasu also served as a great teaching mentor, which made my three years of graduate teaching a rewarding experience. I am lucky to have the empirical and industry-specific wisdom embodied in both Drs. Singhal and Ku, and thank them for their insights and for keeping me on the right path as my dissertation committee members. Even I as recall taking his statistics course, as an MBA student, Dr. Chang has consistently been a beacon of support. Special thanks also to Drs. Ramachandran and Tereyagoglu for their generous and valuable feedback and guidance during my final year.

With invaluable mentors and guides, the journey was made sweeter with good companions. Thanks to all my fellow PhD students, especially those whose kind words and deeds kept me going. Jeremy served as my initial guide, and made sure I was always on the right track during my first year. While Chris, my office mate, provided ongoing encouragement and support, and kept me going through to year 4, when he graduated. I was delighted to meet my new office mate Mayank, who simply made year 5 a blast. Thanks too to the senior OM/KM students Gulru and Wenli for all their helpful assistance. I also thoroughly enjoyed the company and support of my international neighbours across the hall, Vincenzo from Italy and my Caribbean/Latin-American compadre German, from Costa Rica. And to Lori, my birthday twin – thanks.

My family and extended family – you stuck through this with me – bringing meals on wheels and going above and beyond to extend your thoughtfulness and support. Paul,

Betty, Tanya, Mike, Carolyn, John, Elizabeth - there are simply no words. Monique, Gabby, Lisa, Camille, Heather, Ramona, Paula, Zoe, Kim, Kerry, Remi, Stacey, Bruce, Rich – I couldn't have asked for better travel companions. To my KPMG PhD Project family – “Oh the places we went”. For keeping the work/life balance balanced – Madoka, Judith, Anita – I am grateful.

Mom - this is for you - *To the world you might be one person, but to one person you might be the world.*

# TABLE OF CONTENTS

	Page
ACKNOWLEDGEMENTS.....	v
LIST OF TABLES .....	x
LIST OF FIGURES .....	xi
SUMMARY .....	xii
1. INTRODUCTION .....	1
2. THE AMBIDEXTERITY PARADOX: BALANCING THE BENEFITS AND PERILS OF EXPLORATION, EXPLOITATION AND LEARNING FROM FAILURE .....	8
2.1 Introduction.....	8
2.2 Conceptual Framework and Hypotheses.....	14
2.2.1 Model of Induced and Autonomous Learning in the Innovation .....	14
2.2.2 Search and Learning on the NK Landscape.....	16
2.2.3 Impact of Exploration on Innovation Performance Outcomes.....	18
2.2.4 Impact of Exploitation on Innovation Performance Outcomes.....	19
2.2.5 Benefits and Perils of Ambidexterity – Balancing Exploration and Exploitation ..	20
2.2.6 Learning from Prior Success and Prior Failure Experience.....	22
2.2.7 Moderating Impact of Prior Success and Prior Failure Experience on Search .....	25
2.3 Data and Empirical Setting .....	29
2.3.1 Dependent Variables.....	30
2.3.2 Independent Variables.....	31
2.4 Methods .....	33
2.5 Results.....	35
2.5.1 Robustness Tests .....	43
2.6 Implications for Theory and Practice.....	46
2.7 Future Research .....	52
3. A DYNAMIC MODEL OF KNOWLEDGE CREATION FROM EXPLORATION AND EXPLOITATION.....	55
3.1 Introduction.....	55
3.2 Related Literature.....	59
3.2.1 Exploration, Exploitation and Temporal Ambidexterity .....	59



3.2.2	Generating and Resolving Uncertainty in Innovation.....	60
3.2.3	Dynamic Decision Making Under Uncertainty.....	61
3.3	A Dynamic Model of Knowledge Creation from Exploration and Exploitation .....	62
3.4	Manager's Performance Objectives.....	67
3.4.1	Short-term versus Long-term Performance Objectives.....	67
3.4.2	Evaluating Innovation Performance.....	68
3.4.3	Cost Objectives .....	71
3.4.4	Objective Function.....	71
3.5	Optimal Dynamic Knowledge Creation Strategies.....	72
3.5.1	Sequential versus Fixed-Dominant Knowledge Creation Strategies .....	78
3.5.2	Front-Loaded versus Back-Loaded Knowledge Creation Strategies .....	80
3.6	Numerical Analysis and Managerial Insights .....	84
3.7	Conclusions.....	91
4.	A DIFFERENTIAL GAME MODEL OF EXPLORATION AND EXPLOITATION UNDER CO-OPETITION .....	96
4.1	Introduction.....	96
4.2	Related Literature.....	99
4.2.1	Innovation and Competition.....	99
4.2.2	Innovation and Cooperation.....	100
4.2.3	Innovation and Learning under Co-opetition.....	100
4.2.4	Exploration, Exploitation and Risk-Taking in Innovation.....	102
4.3	A Model of Exploration and Exploitation with Cooperation and Competition .....	104
4.3.1	A Dynamic Model of Knowledge Creation .....	104
4.3.2	Two-way Knowledge-Sharing of Exploration and Exploitation Knowledge .....	107
4.3.3	Competing on Knowledge in a Right-tail and Left-tail Race.....	110
4.4	Optimal Solutions .....	113
4.4.1	Case 1: Knowledge Sharing in an Exploration-only Alliance .....	116
4.4.2	Case 2: Knowledge Sharing in an Exploitation-only Alliance .....	128
4.5	Managerial Insights and Conclusions .....	131
	APPENDIX A.....	134
	APPENDIX B .....	144
	APPENDIX C .....	175
	REFERENCES .....	186

## LIST OF TABLES

	Page
Table 3.1: Possible Temporal Ambidexterity Strategies .....	78
Table 3.2: Optimal Temporal Ambidexterity Scenarios.....	85
Table 4.1: Special Cases .....	115
Table A1: Range of Exploration-Exploitation Strategies .....	134
Table A2: Descriptive Statistics and Correlations .....	134
Table A3: Complementary Log-Log Model for Probability of Success and Failure .....	135
Table A4: Linear Regression for Mean and Variance of Citations Received.....	136
Table A5: Quantile Regression for Citations Received.....	137
Table A6: Seemingly Unrelated Regression (SUR) for Citations Received .....	138
Table A7: Complementary Log-Log Model for Probability of Success.....	139
Table A8: Zero Inflated Negative Binomial Model for Citations Received.....	140
Table A9: Complementary Log-Log Model, Sub-sample of Publicly Listed Firms .....	141
Table A10: Complementary Log-Log Model for Probability of Success with Alternate Exploit Measure .....	142
Table A11: Complementary Log-Log Model for Probability of Failure with Alternate Exploit Measure .....	143
Table A12: Complementary Log-Log Model for Probability of Success and Failure with Square Root Discount Factor for Experience .....	144
Table A13: Complementary Log-Log Model for Probability of Success and Failure with Linear Discount Factor for Experience.....	145
Table B1: Model Notation .....	146
Table B2: Numerical Analysis Parameter Settings.....	165
Table C1: Model Notation .....	175
Table C2: Numerical Analysis Parameter Settings.....	183

## LIST OF FIGURES

	Page
Figure 2.1: Conceptual Model of Knowledge Creation in the Innovation Process .....	16
Figure 2.2: Interdependence of Exploration, Exploitation and Prior Experience.....	51
Figure 3.1: Innovation Process Dynamics and Objectives .....	72
Figure 4.1: Case 1A-I.....	122
Figure 4.2: Case 1A-II .....	122
Figure 4.3: Case 1A-III .....	123
Figure 4.4: Case 1B-I for firm $j=1$ .....	127
Figure 4.5: Case 1B-III for firm $j=2$ .....	128
Figure 4.6: Case 1B-II for Firm $j=1$ .....	128
Figure B1: Case I $\lambda_1(T)<0, d\lambda_1(T)/dt>0$ .....	150
Figure B2: Case II $\lambda_1(T)>0, d\lambda_1(T)/dt<0$ .....	153
Figure B3: Case III $\lambda_1(T)<0, d\lambda_1(T)/dt<0$ .....	155
Figure B4: Case IV $\lambda_1(T)>0, d\lambda_1(T)/dt>0$ .....	159
Figure B5: Case 1A Short-Term Risk Averse and Long-Term Risk Averse .....	166
Figure B6: Case 1B Short-Term Risk Averse and Long-Term Risk Averse.....	167
Figure B7: Case 1C Short-Term Risk Averse and Long-Term Risk Averse.....	168
Figure B8: Case 1D Short-Term Risk Averse and Long-Term Risk Averse .....	169
Figure B9: Case 1E Short-Term Risk Averse and Long-Term Risk Averse.....	170
Figure B10: Case 2A Short-Term Risk Seeking and Long-Term Risk Seeking .....	171
Figure B11: Case 2B Short-Term Risk Seeking and Long-Term Risk Seeking.....	172
Figure B12: Case 3 Short Term Risk Seeking and Long Term Risk Averse .....	173
Figure B13: Case 4 Short Term Risk Averse and Long Term Risk Seeking .....	174
Figure C1: Solutions for $\lambda_\sigma(t)$ and $\lambda_s(t)$ under Case 1A.....	179
Figure C2: Possible solutions for $\lambda_\sigma(t)$ .....	180
Figure C3: Solutions for $\lambda_\sigma(t)$ and $\lambda_s(t)$ under Case 1B.....	181
Figure C4: Case 1B-I .....	183
Figure C5: Case 1B-II for firm $j=1$ and Case 1B-III for firm $j=2$ .....	184

## SUMMARY

The discovery and successful development of a technology innovation requires dual capabilities to both explore new knowledge as well as to exploit existing knowledge (March 1991). Innovation, therefore, requires an ambidextrous knowledge creation strategy, defined as the simultaneous pursuit of both exploration and exploitation (O'Reilly and Tushman 2008, Andriopoulos and Lewis 2009, Raisch and Birkinshaw 2008, Lavie et. al 2010, O'Reilly and Tushman 2013). While the benefits of ambidexterity at the firm level are well-accepted, that is, the benefits of diversifying exploration and exploitation across different organizational sub-units or functional domains, there is a call for a more granular task-level examination of the ambidexterity phenomenon (Raisch and Birkinshaw 2008). In response to this call, this thesis considers the challenges of pursuing an ambidextrous knowledge creation strategy. This research generates insights for synchronizing exploration and exploitation activities within an innovation project. It has particular importance for understanding and managing the dynamic evolution and resolution of uncertainty during the innovation process.

A *temporal ambidexterity* strategy is one in which a single organizational unit dynamically balances its investments in exploration and exploitation over time (Brown and Eisenhardt 1997, Siggelkow and Levinthal 2003). All three essays, presented herein, provide new insights on various factors which should be considered when developing and executing a temporal ambidextrous strategy. In the first essay (Chapter 2), I examine the impact of exploration, exploitation and learning from cumulative innovation experience on the likelihood of successfully versus unsuccessfully generating a breakthrough innovation. The first essay provides empirical support for the two analytical essays which

follow. Specifically, I demonstrate three important tenets for developing a theory of temporal ambidexterity. First, I confirm, as conceptually expected (March 1991), that when pursued independently, exploration and exploitation have opposing variance-generating versus variance-reducing impacts on innovation performance, respectively. Second, I show that exploration and exploitation have a negative interaction effect on innovation performance. So that, in the short-term, jointly pursuing exploration and exploitation reduces the likelihood of an innovation breakthrough. Third, I find that the benefits of ambidexterity accrue in the long-term, as a result of learning from prior failure experience. This result provides empirical support for the benefits implied by the innovation management mantra: “fail fast, fail often”. However, I also demonstrate the boundary conditions of learning from failure. Specifically, I demonstrate that prior failure experience and exploitation are jointly necessary, but not independently sufficient, for learning from failure to occur. Furthermore, the empirical results also demonstrate the potential for succumbing to either an “exploration failure trap” or an “exploitation success trap”, with cumulative failure and success experience, respectively. In summary, the results of Essay 1 demonstrate that pursuing an ambidextrous knowledge creation strategy necessitates a delicate balancing act, in order to manage the short-term and long-term benefits, and perils, of exploration, exploitation and prior innovation experience.

In Essay 2 (Chapter 3), I introduce a dynamic optimization model of temporal ambidexterity, which extends the empirical findings from Essay 1 (Chapter 2). I examine the optimal sequencing of exploration and exploitation knowledge creation activities throughout the innovation process. I consider how an innovation manager’s optimal dynamic investments in exploration and exploitation are driven by the innovation team’s

knowledge creation capabilities and prior innovation experience, and by the manager's short-term and long-term innovation risk objectives. The results demonstrate the conditions under which various temporal ambidexterity strategies endogenously arise. Interestingly, I show that the colloquial “fail early, fail cheap” strategy, that is a strategy focused on early variance-generating exploration, is the optimal strategy for a risk-averse manager, when the innovation team's prior experience is such that the team is initially able to generate a predictable, but limited, range of innovation performance outcomes, and given low marginal costs of exploration.

Finally, in Essay 3 (Chapter 4), I extend the single firm model introduced in Essay 2 (Chapter 3), to develop a model of temporal ambidexterity for two firms jointly pursuing knowledge creation and knowledge-sharing under co-opetition. Here, I consider how co-opetition, that is, cooperative knowledge-sharing with a competitor, impacts a firm's optimal ambidextrous knowledge creation strategy. Specifically, I compare the optimal knowledge creation and knowledge-sharing strategies under two competitive regimes: (i) competition to achieve the best relative performance (risk-seeking) and (ii) competition to avoid the worst relative performance (risk-averse). I consider two-way knowledge sharing, and I assume that each firm freely reveals its knowledge to its competitor, without receiving compensation. The dynamic analytical results contribute to the open questions regarding optimal knowledge-sharing strategies under co-opetition by demonstrating not only “*how much*” and “*what knowledge should be shared*” but, also “*when*” and “*under what conditions*” knowledge-sharing with a co-opetitive partner is beneficial (Loebecke et al. 1999). Importantly, I analytically examine the factors which drive empirically observed alliance dysfunctions, wherein organizations delay

knowledge-sharing and withhold information from their alliance partners (Hamel 1991, Khanna et al. 1998, Müller 2010).

# **CHAPTER 1**

## **INTRODUCTION**

This research aims to deepen our understanding of the knowledge creation challenges which occur during the innovation process. The discovery and successful development of a technology innovation requires dual capabilities to both explore new knowledge as well as to exploit existing knowledge (March 1991). Innovation, therefore, requires an ambidextrous knowledge creation strategy, defined as the simultaneous pursuit of both exploration and exploitation (O'Reilly and Tushman 2008, Andriopoulos and Lewis 2009, Lavie et. al 2010, Raisch and Birkinshaw 2008, Lavie et. al 2010, O'Reilly and Tushman 2013). However, exploration and exploitation have opposing variance-generating versus variance-reducing effects on innovation outcomes. As such, they have been characterized as conflicting modes of knowledge creation, which require different skills and processes, and which are difficult to pursue simultaneously, in the same space and/or time. At the same time, at the firm level, exploration and exploitation have been shown to have complementary effects (He and Wong 2004), which suggests that they should be pursued together in some optimal balance (March 1991).

While the benefits of ambidexterity at the firm level are well-accepted (Lavie et. al 2010, Raisch and Birkinshaw 2008, Lavie et. al 2010, O'Reilly and Tushman 2013), that is, the benefits of diversifying exploration and exploitation across different organizational sub-units or functional domains, there is a call for a more granular task-level examination of the ambidexterity phenomenon (Raisch and Birkinshaw 2008). In response to this call, this thesis considers the challenges of pursuing an ambidextrous



knowledge strategy within a single innovation project. As opposed to an organizational ambidexterity strategy, which balances exploration and exploitation across various organizational units or functional domains, a *temporal ambidexterity* strategy is one in which a single organizational unit dynamically balances its investments in exploration and exploitation over time (Brown and Eisenhardt 1997, Siggelkow and Levinthal 2003). All three essays, presented herein, provide new insights on various factors which should be considered when developing and executing a temporal ambidexterity strategy.

The first essay (Chapter 2) is an empirical study that examines the impact of pursuing an ambidextrous knowledge creation strategy on innovation performance outcomes. Based on a sample of patents, granted for innovations in the biomedical device industry, I examine how firms can effectively synchronize exploration and exploitation, in order to maximize the upside potential for innovation success while minimizing the downside risk of failures. For the purposes of this study, I refer to an innovation as a success if the innovation process results in the creation of a breakthrough innovation, which is evaluated as having above-average technological value (i.e., an extreme right-tail realization), relative to a population of comparable innovations. Conversely, I define an innovation outcome as a failure when the innovation is evaluated as having below-average technological value (i.e., an extreme left-tail realization), relative to a population of comparable innovations.

Two competing views exist on the most effective strategies for improving the likelihood of success from exploration, while mitigating the associated high levels of uncertainty and failure. One view advocates pursuing an ambidextrous strategy. As an alternative, other researchers suggest that, instead of simultaneously pursuing

exploitation in order to reduce risks of exploration, managers should actively embrace the uncertainty associated with exploration. Researchers and practitioners argue that, while exploration may lead to failure in the short-term, failures improve the likelihood of generating breakthrough innovations in the long-term. Consistent with this view, innovation incentive structures and funding systems which embrace uncertainty and exhibit a high tolerance for failure are encouraged (Tian and Wang 2011, Azoulay et al. 2011). This notion, regarding the benefits of failure, is promoted in the often-cited innovation mantra: “fail fast, fail often” (Thomke 2001). Yet, according to Cannon and Edmonson (2001): “Despite the importance of learning from failure, however, it is more common in exhortation than in practice, and our understanding of the conditions under which it occurs is limited.” (p.161).

The results of Essay 1 (Chapter 2) provide a critical link between theories of ambidexterity and theories of learning from failure. I show that exploration and exploitation improve innovation performance through two separate pathways: (i) exploration directly increases the likelihood of a successfully generating a breakthrough innovation and (ii) exploitation and prior failure experience jointly increase the likelihood of generating a breakthrough innovation. Therefore, I demonstrate the critical role of exploitation for enabling learning from prior failure experience. These results suggest that, rather than being alternate strategies for improving innovation performance, in fact, ambidexterity and learning from failure are complementary processes which operate in tandem.

I also demonstrate other features salient to understanding the paradoxical challenges of ambidexterity. First, I demonstrate that exploration increases the variance

of innovation outcomes, as well as the likelihood of generating a breakthrough. The opposite is true for exploitation. Second, I demonstrate a negative interaction effect between exploration and exploitation, which highlights the challenge of pursuing ambidexterity. Third, I demonstrate that prior success experience can lead to an “exploitation success trap”. Finally, I find that prior failure experience can lead to an “exploration failure trap”. Collectively, the results illustrate that pursuing an ambidextrous knowledge creation strategy necessitates a delicate balancing act, in order to manage the short-term and long-term benefits, and perils, of exploration, exploitation, and learning from failure.

In Essay 2 (Chapter 3), I introduce a dynamic optimization model of temporal ambidexterity, which extends the empirical findings from Essay 1 (Chapter 2). The model considers the manager of an innovation team who invests in exploration and exploitation during the innovation process. The uncertain nature of the development activities is captured in terms of the mean and variance of the distribution of possible innovation performance outcomes. The manager must determine how to balance the variance-generating effects of exploration against the variance-reducing effects of exploitation. A key feature of the dynamic model is the consideration of the innovation team’s absorptive capacity (Cohen and Levinthal 1990), which is a function of its prior innovation experience. Absorptive capacity can both improve as well as constrain the path for future learning. I assume that past exploration generates future opportunities for exploitation (Rothaermel and Deeds 2004, Lavie et al. 2010). On the other hand, past exploitation tends to reduce opportunities for further exploitation (Fleming 2001). Furthermore, I assume a lagged realization of the performance benefits of exploration. Importantly, I

consider how the manager's decision to invest in exploration and exploitation is impacted by his short-term versus long-term risk objectives. These dynamic risk preferences are a function of the level of technical uncertainty which remains unresolved at a given point in time during the innovation process. I examine the optimal sequencing of knowledge creation activities, and provide examples where either the typical explore-then-exploit sequential strategy or the atypical exploit-then-explore sequential strategy is optimal. Interestingly, I show that the colloquial "fail early, fail cheap" strategy, that is a strategy focused on early variance-generating exploration, is the optimal strategy for a risk-averse manager, when the innovation team's prior experience is such that the team is initially able to generate a predictable, but limited, range of innovation performance outcomes, and given low marginal costs of exploration.

In Essay 3, (Chapter 4) I extend the model of temporal ambidexterity, introduced in Essay 2 (Chapter 3), to introduce a differential game model of knowledge-sharing between two rival firms. In the third essay, I consider how co-opetition, that is, cooperative knowledge-sharing with a competitor, impacts a firm's optimal ambidextrous knowledge creation strategy. A firm can participate in knowledge-sharing alliances in order to explore new technological opportunities, as well as to improve the ability to exploit its existing capabilities. Therefore, I consider two alternative types of alliances: (i) an exploration knowledge-sharing alliance and (ii) an exploitation knowledge-sharing alliance. Based on March's (1991) framework of competition for relative position in a right-tail race versus in a left-tail race, I compare the optimal knowledge creation and knowledge-sharing strategies under two different competitive performance regimes:

(i) competition to achieve the best relative performance (risk-seeking) and (ii) competition to avoid the worst relative performance (risk-averse).

I examine the optimal sequencing of exploration and exploitation activities, under both types of alliances, and provide examples where either the typical explore-then-exploit sequential strategy or the atypical exploit-then-explore sequential strategy may be optimal under co-opetition. For both the exploration and exploitation knowledge-sharing alliance, I consider two-way knowledge sharing in which I assume that each firm freely reveals its knowledge to its competitor, without receiving compensation. As a result of this modeling assumption, I am able to gain a better understanding of a firm's incentives for free-revealing and participating in cooperative innovation with its competitor (Von Hippel and von Krogh 2003, 2006). Importantly, the analytical results provide insights on how a firm should optimally manage its knowledge sharing exchanges, in order to balance the potential benefits of cooperation, against the potential threat of being "out-learned and out-competed by the competitor-partner" (Gnyawali and Park 2011, p. 657). The results also provide a better understanding of the motivating factors which drive empirically observed alliance dysfunctions, wherein organizations delay knowledge-sharing and withhold information from their alliance partners (Hamel 1991, Khanna et al. 1998, Müller 2010).

Collectively, all three essays provide insights for managers tasked with allocating knowledge creation resources during the innovation process (Gaimon and Bailey 2012). Specifically, this research generates insights for synchronizing exploration and exploitation activities within an innovation project. It has particular importance for understanding and managing the dynamic evolution and resolution of uncertainty during

the innovation process. The findings presented herein will enable innovation managers to more effectively leverage the benefits of an ambidextrous knowledge creation strategy, by taking into consideration issues related to risk-management, short-term and long-term performance tradeoffs, learning from prior innovation experience and leveraging external knowledge sources under competition. An innovation process in which exploration precedes exploitation has been accepted as the typical sequence of knowledge creation activities, however, the findings of these three studies suggest conditions under which alternate strategies may be optimal. This suggests that a more adaptive approach is required when synchronizing exploration and exploitation knowledge creation activities during the innovation process.

## **CHAPTER 2**

# **THE AMBIDEXTERITY PARADOX: BALANCING THE BENEFITS AND PERILS OF EXPLORATION, EXPLOITATION AND LEARNING FROM FAILURE**

### **2.1 Introduction**

Firms which are able to successfully undertake the process of innovation can reap benefits in terms of competitive advantage (Barney 1991), increased financial returns (Sorescu et al. 2003) and improved market value (Kelm et al. 1995, Sood and Tellis 2009). To demonstrate the potential returns from successful innovation, consider the global market for medical devices which was estimated to be worth more than \$200 billion dollars at a mid-2000 estimate (Denend and Zenios 2006). Several high profile examples illustrate the financial benefits which result from innovation breakthroughs within the medical devices industry. For example, Boston Scientific Corporation's access to license the patent for a breakthrough innovation for cardiac stents, branded Taxus, was estimated to potentially double the medical device firm's \$3.5 billion annual revenues (Kerber 2004). Similarly, in 2011 Medtronic acquired medical device developer Ardian for \$800 million, namely to gain access to its breakthrough renal denervation system branded Simplicity (Businesswire 2010). While in 2013, IDEV Technologies, which develops next generation medical devices, was acquired by Abbott for \$310 million primarily to gain proprietary access to its stent system branded Supera Veritas (Abbott Laboratories 2013).

Given their potential to generate significant firm value, researchers have sought to better understand the sources of breakthrough innovations. Hereafter, we refer to an innovation as a *success* if the innovation process results in the creation of a breakthrough innovation, which is evaluated as having above-average technological value (i.e., an extreme right-tail realization), relative to a population of comparable innovations (Anderson and Tushman 1990). Conversely, we define an innovation outcome as a *failure* when the innovation is evaluated as having below-average technological value (i.e., an extreme left-tail realization), relative to a population of comparable innovations. The degree to which a firm chooses to focus on exploration, that is, broad search and recombination across multiple knowledge domains to generate new knowledge, versus exploitation, that is building on existing knowledge domains (March 1991), has been identified as a critical predictor of inter-firm differences in innovation performance (He and Wong 2004, Rothaermel and Deeds 2004). However, March (1991) cautions against the dangers of pursuing either mode of knowledge creation in excess.

While exploration, that is the search for new knowledge across multiple, knowledge domains, has been empirically shown to be positively associated with generating breakthrough innovations (Rosenkopf and Nerkar 2001, Katila and Ahuja 2002), an excessive emphasis on exploration is also expected to increase uncertainty and the variance of the distribution of innovation outcomes (March 1991, Fleming and Sorenson 2001, He and Wong 2004). As a result, while exploration increases the variance and the likelihood of a breakthrough innovation (i.e., the likelihood of realizing an extreme right tail outcome) the increased variance is also expected to increase the likelihood of failure (i.e., the likelihood of realizing an extreme left tail outcome). On the



other hand, a focus on exploitation of existing knowledge is expected to decrease uncertainty and reduce the risk of failure (March 1991). However, an excessive focus on exploitation of existing knowledge leads to generating only incremental innovation outcomes, with diminishing returns (Sørensen and Stuart 2000, Ahuja and Lampert 2001).

Since exploration is accepted as a core requirement for generating breakthrough innovations, an obvious question arises: What is the optimal way for innovators to mitigate the high levels of risk and uncertainty associated with exploration? Two competing views exist on the most effective strategies for improving the likelihood of breakthrough success from exploration, while also mitigating the high levels of uncertainty and failure associated with exploration. One view advocates pursuing an ambidextrous strategy, that is, advocates simultaneously pursuing both exploration and exploitation (Andriopoulos and Lewis 2009, Raisch and Birkinshaw 2008, Lavie et. al 2010, O'Reilly and Tushman 2013). An ambidextrous knowledge creation strategy is purported to allow innovators to achieve the best of both worlds, by balancing the variance-increasing benefits of increased breakthroughs from exploration with the variance-reducing benefits of exploitation (He and Wong 2004). Contrary to the ambidexterity view, other researchers and practitioners suggest that, instead of simultaneously pursuing exploitation in order to reduce risk and uncertainty of exploration, innovators should instead actively embrace the increased uncertainty associated with exploration. These researchers argue that, while exploration may lead to failure in the short-term, failures provide innovators with the opportunity to improve the likelihood of generating successful breakthroughs in the long-term (Cannon and

Edmonson 2005). Consistent with this notion, recent research has shown that organizations with incentive structures and funding systems which embrace uncertainty and exhibit a higher tolerance for failure yield higher rates of innovation (Tian and Wang 2011, Azoulay et al. 2011). This view regarding the benefits derived from embracing failure is particularly widespread, and is promoted in the often-cited innovation mantra, “fail fast, fast often” (Thomke 2001). Nevertheless, the value of failure is not yet well-understood. According to Cannon and Edmonson (2001): “Despite the importance of learning from failure, however, it is more common in exhortation than in practice, and our understanding of the conditions under which it occurs is limited.”

In this paper, we introduce and empirically test a model of knowledge creation and ambidexterity and examine questions related to the two competing views discussed above. To examine the role of ambidexterity we pose: *Do the benefits derived from pursuing an ambidextrous knowledge creation strategy outweigh the benefits of independently pursuing exploration? What role does exploitation play in an ambidextrous knowledge creation strategy?* To examine the role of learning we ask: *Does learning from prior successes and/or learning from prior failures improve the likelihood of future successes?* Employing a sample of patents in the biomedical devices industry, we investigate the impact of exploration, exploitation, prior success experience and prior failure experience on future innovation outcomes. Our results highlight the challenges and benefits of simultaneously pursuing exploration and exploitation, in the short-term, as well as the role of learning from success and learning from failure, for realizing the long-term benefits of ambidexterity.

Our findings contribute to the literature in several ways. First, we confirm that exploration increases the likelihood of generating a breakthrough. Furthermore, as conceptually expected, we also empirically demonstrate that exploration increases the variance (uncertainty) of innovation outcomes. Secondly, while we find that exploitation reduces the variance (uncertainty) of innovation outcomes, we also demonstrate that exploitation has negative implications for innovation performance, because it reduces the likelihood of breakthrough successes. Furthermore, contrary to the studies which demonstrate a positive benefit to ambidexterity, at the firm level and at later stages in the innovation process, (Katila and Ahuja 2002, He and Wong 2004, Rothaermel and Deeds 2004, Cao et al. 2009, Chandrasekaran et al. 2011), our results, which focus on the initial knowledge creation stage of the innovation process, indicate a negative interaction between exploration with exploitation. Moreover, contrary to the proposition that innovators learn from failure, our results suggest that innovators learn more readily from prior success experience, since we find that prior failure experience has a direct negative impact on innovation performance. Interestingly, we demonstrate, through a positive interaction term, that prior failure experience along with exploitation are jointly necessary, but are not independently sufficient, for learning from failure to occur. In essence, we uncover the benefits, as well as the limitations, of exploitation and learning from failure.

Our findings also provide a critical link between theories of ambidexterity and theories of learning from failure in the innovation process. We show that exploration and exploitation improve innovation performance through two separate pathways: (i) exploration directly increases the likelihood of a success outcome and (ii) exploitation

and prior failure experience are jointly necessary, but not independently sufficient, to increase the likelihood of a future success outcome. Taken together, the results indicate that innovation performance is enhanced by undertaking both exploration and exploitation, that is, pursuing ambidexterity, however, the benefits of exploitation accrue from the ability to effectively learn from prior failure experience. These results suggest that, rather than being alternate strategies for improving innovation performance, in fact, ambidexterity and learning from failure are to be seen as complementary processes which operate in tandem with each other. Our results on the benefits of prior failure experience also provide support for the veracity of the innovation mantra: “fail fast, fail often” (Thomke 2001).

Our finding on the role of prior success and failure experience, as a moderator in the knowledge creation and innovation process, is also an important contribution. Although prior research has analytically modeled the role of learning and memory in the innovation search process (Jain and Kogut 2013), the impact of learning from experience as a moderator in the search process has not been empirically tested. Furthermore, to our knowledge, no prior research has simultaneously compared the impact of prior success experience and prior failure experience on the generation of subsequent innovation breakthroughs. Collectively, our empirical findings and insights contribute to a deeper conceptual understanding of knowledge creation, learning and ambidexterity in the innovation process.

## **2.2 Conceptual Framework and Hypotheses**

### **2.2.1 Model of Induced and Autonomous Learning in the Innovation**

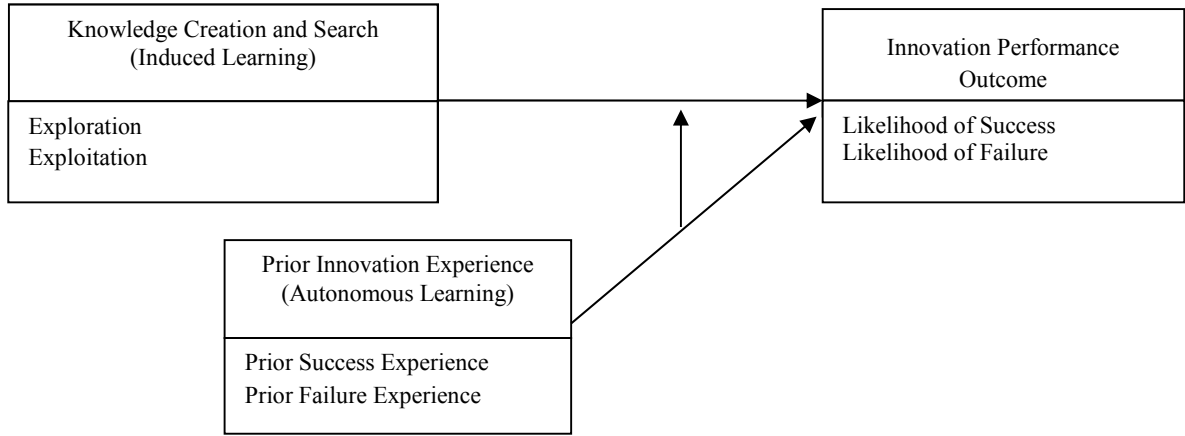
We propose a model of the impact of an innovator's knowledge creation activities on the distribution of innovation outcomes. In developing this model, we build on the literature related to organizational learning which suggests that induced learning and autonomous learning can serve as substitutes and complements for improving performance (Adler and Clark 1991, Ittner et al. 2001, West and Iansiti 2003). Our model is unique in that we consider the interdependence between induced and autonomous learning within the context of the innovation search process. Whereas autonomous learning refers to learning-by-doing and recognizes cumulative experience as a source of knowledge which contributes to future performance improvements (Yelle 1979, Argote 1991), induced learning refers to the deliberate activities and investments which are made to improve performance. In our model, we consider two alternative induced learning activities in the innovation search process: exploration and exploitation (March 1991). While the exploration/exploitation tradeoff can be conceptualized as two ends of a continuum (Gupta et al. 2006), we follow Katila and Ahuja's (2002) definition of exploration and exploitation as orthogonal. That is, we assume innovators choose to invest in varying degrees of exploration as well as varying degrees of exploitation during the innovation process, so that they can choose to operate in one of four quadrants within this two-dimensional framework. Our choice of representing exploration and exploitation as orthogonal measures allows us the flexibility to test two key elements of our model: (i) the interaction between exploration and exploitation, as a measure of ambidexterity and (ii) the moderating effect of prior innovation experience on the effectiveness of both

exploration and exploitation in the knowledge creation and innovation process. We introduce our conceptual model below.

While building cumulative experience, innovators observe both success and failure outcomes as a result of previous activities. Based on these outcomes, innovators are able to make inferences, which enable them to better predict cause-and-effect, in order to improve future performance. The recognition that success and failure outcomes each hold idiosyncratic information, has ignited a stream of research in the areas of learning from failure and learning from success (Kim et al. 2009, Madsen and Desai 2010, KC et al. 2013). While most studies have focused on the benefits of either learning from failure or learning from success, recent studies have begun to directly compare the benefits of learning from failure versus learning from success (Madsen and Desai 2010). We extend these recent studies which consider success experience and failure experience as dual inputs in the innovation process.

Broadly speaking, the performance benefits of learning are defined as either (i) achieving superior performance, or (ii) reducing performance variability, failures or errors. Lavie et al. (2010) suggest that an investigation of exploration and exploitation necessarily requires a dual outcome measure which captures the inherent performance tradeoffs between these two modes of knowledge creation. In this paper we consider the impacts of exploration, exploitation and learning from experience on the likelihood of success versus failure, in the knowledge creation and innovation process. In doing so, we address the need for a two-fold measure which can quantify these tradeoffs between exploration and exploitation in maximizing the upside potential for generating breakthroughs versus minimizing the downside risk of failure. In summary, the model of

knowledge creation examines how an innovator’s exploration and exploitation activities (induced learning) as well as prior success experience and prior failure experience (autonomous learning) act as both substitutes and complements for improving innovation performance (Adler and Clark 1991, Ittner et al. 2001, West and Iansiti 2003) (See Figure 2.1).



**Figure 2.1: Conceptual Model of Knowledge Creation in the Innovation Process**

## 2.2.2 Search and Learning on the NK Landscape

Following Fleming (2001) we consider innovation as a process of recombinant search. As such, we assume that innovations result when innovators recombine various technological components in new ways. We leverage the analogy of the *NK* performance landscape in order to demonstrate the impact of an innovator’s search activities over a rugged landscape of uncertain innovation performance outcomes (Levinthal 1997, Fleming and Sorenson 2001, Fleming and Sorenson 2004, Erat and Kavadias 2008, Jain and Kogut 2013). Within the *NK* model, *N* refers to the number of technological domains over which the innovator’s recombination effort is applied, and *K* refers to the degree of interdependence between these *N* domains, as it relates to the value of the innovation output. As the number of technological domains *N* and the degree of interdependence *K*

rises, the landscape of performance becomes increasingly rugged; that is the innovation performance landscape has increasingly higher peaks as well as lower valleys, with respect to the realized value of the innovation output. The peaks (valleys) in the performance landscape correspond to different technology configurations which yield superior (inferior) performance. However, a-priori, these configurations are unknown to the innovator.

The literature recognizes two different strategies for traversing the  $NK$  landscape. The process of exploration, or “long jumps”, refers to searching for new innovation outcomes by recombining multiple components (high  $N$ ), across distant or unrelated domains (high  $K$ ) (Levinthal 1997). However, given the vast and rugged expanse of the innovation performance landscape, as well as the limited knowledge of the innovator, exploration effectively becomes a process of “blind search” over the landscape. Alternatively, therefore, innovators may leverage a “map” of the landscape, which they use as a guide in their search efforts (Fleming and Sorenson 2004). The map of the landscape reflects existing scientific research and provides a codified relationship of cause-and-effect, which improves the innovators ability to make predictions between moves to various positions of the landscape and the resulting performance outcomes (Fleming and Sorenson 2004). Such reliance on extant scientific research in the innovation search process, is consistent with March’s (1991) definition of exploitation as “refinement of an existing technology” (p. 72). However, Levinthal (1997) notes that the outcome of the innovation search process may also be influenced by the starting point on the innovation performance landscape. This highlights the importance of considering the role of prior innovation experience. As innovators repetitively attempt to generate novel



breakthroughs within an unknown rugged landscape they generate knowledge by trial-and-error learning (Nelson 2008, Rerup and Feldman 2011). Furthermore, this prior innovation experience necessarily modifies the effectiveness and efficiency of future search efforts (Cohen and Levinthal 1991). Naturally, the notion of trial-and-error learning suggests that the result from each innovation attempt could either be a successful trial outcome or a failed trial outcome. However, the explicit role of success experience versus failure experience as a moderating factor in the innovation search process has not been hypothesized or empirically tested. Therefore, in this paper, we consider how learning from prior success and prior failure experience critically influences the effectiveness and efficiency of future exploration and exploitation search activities across the rugged *NK* landscape.

### **2.2.3 Impact of Exploration on Innovation Performance Outcomes**

Based on the *NK* model of search, the process of exploration, or “long jumps”, refers to recombining multiple components (high *N*), across distant or unrelated domains (high *K*). Managerially, exploration is associated with increasing the scope of search (Katila and Ahuja 2002), increasing the diversity and breadth of knowledge sources (Taylor and Greve 2006, Lazer and Friedman 2007, Wu and Shanley 2009), and accessing external knowledge sources (Rosenkopf and Nerkar 2001). This strategy of combining a large number of seemingly unrelated technological components underscores the principles of creativity and brainstorming, which give rise to successful innovation outcomes, by scaling to one of the peaks on the rugged landscape. However, as the innovation performance landscape becomes more rugged, outcomes from exploration becomes more uncertain and unpredictable (Fleming and Sorenson 2001, Austin et al.

2012). As a result of the increased variance of the outcomes, the search may either result in generating superior outcomes (scaling a peak), or undesirable outcomes (traversing a valley). Therefore, we posit that increasing the degree of exploration increases the potential to generate a success, but also increases the likelihood of realizing a failure in the innovation process (Austin et al. 2012).

***HYPOTHESIS 1: Exploration increases the likelihood of both success (extreme right-tail outcome) and failure (extreme left-tail outcome) in the innovation process.***

#### **2.2.4 Impact of Exploitation on Innovation Performance Outcomes**

In addition to the degree of exploration, the innovator can also choose the degree of exploitation to pursue for an innovation project. March (1991) describes exploitation as the “refinement of an existing technology”. Therefore, we consider exploitation as being associated with building upon a technological trajectory (Dosi 1982, Benner and Tushman 2003) through the reuse of existing knowledge (Fleming 2001). Innovators build on existing knowledge by relying on available scientific knowledge. Fleming and Sorenson (2004) suggest that reliance on extant scientific knowledge can improve the innovation search process by serving as a “map” of the uncertain performance landscape. As technological knowledge becomes more mature, standardized, codified and well-understood over time, then the “map” becomes more representative of the landscape and increases in predictive power (Hamilton and Singh 1992, Bohn 1994, Benner and Tushman 2003, Autio et al. 2004, Linton and Walsh 2004). Therefore, as exploitation and hence the degree of reliance on the “map” increases, outcomes become more predictable, and uncertainty is reduced. However, the exploitation of existing, mature technologies reduces the potential for generating a novel breakthrough outcome (Sorenson and Stuart

2000, Ahuja and Lampert 2001, Katila and Ahuja 2002, Fleming and Sorenson 2004). In other words, pursuing exploitation entails a performance tradeoff so that “preventing failure can mean sacrificing opportunity” (McGrath 1999). Sorensen and Stuart (2000) establish a related tradeoff by empirically demonstrating that while a focus on exploiting mature technology improves innovation efficiency, the gains are realized at the expense of generating only incremental innovation outcomes. Consistent with these arguments we propose the following hypothesis.

***HYPOTHESIS 2: Exploitation decreases the likelihood of both success (extreme right-tail outcome) and failure (extreme left-tail outcome) in the innovation process.***

### **2.2.5 Benefits and Perils of Ambidexterity – Balancing Exploration and Exploitation**

In his seminal work on exploration and exploitation, March (1991) cautions against the danger of focusing exclusively on either exploration or exploitation, and suggests that a balance between the two modes of knowledge creation may be more appropriate. Several studies confirm the positive benefits, with respect to firm performance, of employing an ambidextrous organizational strategy (Katila and Ahuja 2002, He and Wong 2004, Rothaermel and Deeds 2004, Chandrasekaran et al. 2012). For example, He and Wong (2004) find a positive interaction between exploration and exploitation activities, leading to an increased rate of sales growth. Rothaermel and Deeds (2004) also find that performance improves if firms employ an ambidextrous alliance strategy. While many of the studies of ambidexterity have focused on this phenomenon at the firm level, Raisch and Birkinshaw (2008) underscore the need for a more granular task-level examination of the ambidexterity phenomenon.

Our definitions distinguishing between exploration and exploitation are consistent with several factors considered complementary in the knowledge creation process including: building knowledge breadth versus knowledge depth (McKee 1992, Katila and Ahuja 2002, Wu and Shanley 2009), increasing knowledge diversity versus specialization (Narayanan et al. 2009), focusing on cross-functional versus within-function knowledge (Bajaj et al. 2004), and using novel versus standardized processes (Gilson et al. 2005), respectively. For example Katila and Ahuja (2002) find a positive interaction between knowledge breadth (exploration) and knowledge depth (exploitation) on the number of new products launched, for firms which employ an ambidextrous innovation search strategy. Relatedly, Gilson et.al (2005) find support for the complementary effect between creativity (exploration) and standardization of work processes (exploitation) on team performance in a knowledge-intensive service setting. Therefore in Hypothesis 3A, we posit that pursuing ambidexterity improves innovation performance. We characterize a performance benefit as either an improved likelihood of generating a success, or a reduction in the likelihood of generating failures.

***HYPOTHESIS 3A: Jointly pursuing exploration and exploitation increases the likelihood of success (extreme right-tail outcome) and decreases the likelihood of failure (extreme left-tail outcome) in the innovation process.***

Contrary to the benefits proposed in Hypothesis 3A, the notion of the “ambidexterity paradox” suggests that, although beneficial, an ambidextrous approach may be difficult to implement (Leonard-Barton 1992, Gupta et al. 2006, Raisch and Birkinshaw 2008, Andriopoulos and Lewis 2009). This difficulty results due to the tensions and conflicts which arise when trying to simultaneously undertake the dissimilar

and contradictory processes of exploration and exploitation. For example, exploitation capabilities are associated with the standardization of processes and the codification of existing knowledge. On the other hand, exploration is constrained by the inertia and core rigidities which can result from the same process standardization and knowledge codification activities which enhance exploitation (Leonard-Barton 1992, Benner and Tushman 2003). For example, Gilson et al. (2005) find that an organizational focus on standardization and reliance on established practices enhance exploitative learning, but inhibit exploration and creativity. Relatedly, in a field study, Wong (2004) finds that the organizational factors which enhance external exploratory learning, also inhibit local exploitation learning, and shows a negative interaction between these two activities. Therefore in Hypothesis 3B, we put forward the alternate hypothesis that ambidexterity has negative performance implications.

***HYPOTHESIS 3B:*** *Jointly pursuing exploration and exploitation decreases the likelihood of success (extreme right-tail outcome) and increases the likelihood of failure (extreme left-tail outcome) in the innovation process.*

#### **2.2.6 Learning from Prior Success and Prior Failure Experience**

We also consider how prior innovation experience serves as a starting point which may critically influence the efficiency of future innovation efforts (Cohen and Levinthal 1990, Levinthal 1997). Jain (2013) empirically shows that the cost to complete an innovation is reduced with cumulative experience, suggesting that innovators can in-fact “learn to innovate”. We extend Jain (2013) by distinguishing between learning from success experience versus learning from failure experience. Below we develop

hypotheses regarding the impact of prior success experience and prior failure experience on subsequent innovation performance.

### ***Learning from Success***

While some studies have explicitly examined the impact of learning from success experience (Kim et. al 2009, KC et al. 2013), other studies have implicitly assumed that firms are, in fact, learning from “successful” experiences (Lapre and Nembhard 2010). The management philosophy of implementing “best practices” is a classic example of learning from success (Tucker et al. 2007). A best practice is a solution known to yield a successful result. Identification of prior success factors and best practices can help an innovator to improve performance, by providing a clear set of guiding principles which can be followed to ensure future success (Lee and Van de Steen 2010). Within the context of the search for superior innovations on a rugged landscape, identifying best practices relate to identifying prior technology configurations which have consistently yielded superior outcomes. By observing repeated patterns of prior successful technology configurations, the innovator further refines their “best practices”. Some have referred to this cumulative “best practices” experience, as cognition (Gavetti and Levinthal 2000), memory (Jain and Kogut 2013), templates (Baron and Ensley 2006) or routines (Obstfeld 2012). By building a repository of innovation best practices and thereby improving innovation capabilities, through repeated practice, innovators are able to improve future innovation performance. Therefore in Hypothesis 4, we posit that prior success experience improves innovation performance.

***HYPOTHESIS 4: Prior success experience increases the likelihood of success (extreme right-tail outcome) and decreases the likelihood of failure (extreme left-tail outcome).***

### ***Learning from Failure***

While theories of ambidexterity propose balancing the risks of failure from exploration with the risk-reducing benefits of exploitation, other scholars suggest that failure should be encouraged instead of avoided (McGrath 1999, Thomke 2001, Cannon and Edmonson 2005). Managerial support for such initiatives is reflected in the practice of organizations which provide incentive structures which encourage risk-taking and exhibit a high tolerance for failure. For example, the medical research organization Howard Hughes Medical Institute encourages its researchers to “embrace the unknown even if it means uncertainty or the chance of failure” (Howard Hughes Medical Institute Annual Report 2003, pg. 12). Several well-known examples demonstrate how prior failure can lead to improved likelihood of future success (McGrath 1999). In particular, a failure may suggest alternative outcomes which were not previously considered. For example, a failed experiment at DuPont is credited for the eventual successful discovery of Nylon (Cannon and Edmonson 2005). Furthermore, in addition to helping to generate superior performance outcomes, prior failure experience may reduce the likelihood of subsequent failures. Prior failure can help a firm to reduce the likelihood of future failures, by signifying unsuccessful practices which should be avoided in the future. In support of this view, Haunschild and Sullivan (2002) demonstrate that prior accident experience reduces the likelihood of future accidents. Maden and Desai (2010) also show that failure experience reduces the likelihood of launch failures, in the orbital launch vehicle industry. Therefore in Hypothesis 5, we posit that prior failure experience improves innovation performance.

***HYPOTHESIS 5A:*** *Prior failure experience increases the likelihood of success (extreme right-tail outcome) and decreases the likelihood of failure (extreme left-tail outcome).*

Although the notion of learning from failure has gained considerable attention, some empirical results suggest that the failure experience may not provide performance benefits. Considering innovation as a process of search on a rugged landscape, repeated failures suggest the possibility that an innovator has become stuck or trapped in a valley on the *NK* performance landscape (Sorenson et al. 2006). Furthermore, leveraging theories of organizational behavior, KC et al. (2013) suggest that, based on attribution theory, it may be easier for firms to learn from their success versus their failure experience. Baumard and Starbuck (2005) conduct a number of case studies and find that firms have difficulty leveraging failure experience to improve performance. Moreover, they suggest that repeated failure experience can give rise to maladaptive behaviors, such as an escalating commitment to failure. Therefore, these findings lead us to the alternate hypothesis posited in 5B, below.

***HYPOTHESIS 5B:*** *Prior failure experience decreases the likelihood of generating a success (extreme right-tail) and increase the likelihood of generating a failure (extreme left-tail) outcome.*

### **2.2.7 Moderating Impact of Prior Success and Prior Failure Experience on Search**

Consistent with the notion that failure experience may not provide the expected performance benefits, Sitkin (1992) cautions that, “simply experiencing a negative event is not sufficient for learning” (Cohen and Sproull 1995, p. 142). Moreover, Sitkin (1992) advises that in order to generate learning, intelligent failures must take place in “domains that are familiar enough to permit effective learning” (Cohen and Sproull 1995, p. 145).



This suggests that in a well-understood scientific domain, given a finite set of cause-and-effect relationships, failure experience provides informational value by revealing information about the underlying factors which contributed to the failure (Petkova 2009). This conceptual notion is consistent with Lee and Van de Steen's (2010) results which analytically demonstrate that failure experience has no informational value unless there are a finite number of known solutions.

Pisano (1994) empirically finds that experimental experience only improves performance when there is a sound theoretical knowledge of the science underlying innovation in a particular domain. Nelson (2008) also outlines the role of theory in the problem-solving process required for effective trial-and error learning. Relatedly, Fleming and Sorenson (2004) suggest that exploiting well understood scientific domains may provide benefits in the face of negative feedback; however, they do not empirically test this relationship between exploiting existing technology and prior failure experience. Cannon and Edmonson (2005) provide a related illustration in which Eli Lilly commissioned a mathematician to probe and analyze the underlying causes of failure of the clinical trials for Alimta. By identifying and eventually resolving the underlying sources of failure, the company was able to transform the initial failure into an eventual success. Reinforcing the importance of domain expertise as a prerequisite for learning from failure, in a number of case studies, Baumard and Starbuck (2005) find that "managers did not report on problems that seemed 'out of box' and difficult to explain" (p.293). Consistent with the above arguments, in Hypothesis 6 below, we propose that prior failure experience and exploitation jointly enhance innovation performance.

***HYPOTHESIS 6:*** *The interaction between exploitation and prior failure experience increases the likelihood of generating a success (extreme right-tail) and decreases the likelihood of generating a failure (extreme left-tail) outcome.*

Contrary to the joint benefits of prior failure and exploitation proposed in Hypothesis 6, Levinthal and March (1993) suggest that repeated failure experience can cause innovators to succumb to an “exploration failure trap”. Levinthal and March (1993) describe this potential reinforcing cycle between the exploration of new knowledge and failure as follows: “Failure leads to search and change which leads to new failure which leads to more search and so on. New ideas and technologies fail and are replaced by other new ideas and technologies, which fail in turn.” (p. 106). This suggests that prior failure experience leads to deteriorating innovation performance from exploration, which leads to Hypothesis 7 below:

***HYPOTHESIS 7:*** *The interaction between exploration and prior failure experience decreases the likelihood of generating a success (extreme right-tail) and increases the likelihood of generating a failure (extreme left-tail) outcome.*

On the other hand, Levinthal and March (1993) suggest that repeated success experience can cause innovators to succumb to an “exploitation success trap”. As the organization accumulates success experience, this increases confidence in the current available solutions (Levinthal and March 1993, March 2003). Given that Hypothesis 2 suggests that exploitation decreases the likelihood of both failure and success outcomes, we predict that given prior success experience, these effects are magnified. Therefore, in Hypothesis 8 we predict that prior success moderates the main effect of exploitation, leading to increasingly incremental outcomes, so that that the likelihood of both success

and failure are further reduced. These predictions are consistent with the theory of the “exploitation success trap” (Levinthal and March 1993), which suggests that prior success experience and exploitation jointly reinforce stable, but sub-optimal performance outcomes from the innovation process (Levinthal and March 1993, Leonard-Barton 1992, Audia and Goncalo 2007). Thus, we posit:

***HYPOTHESIS 8:*** *The interaction between exploitation and prior success experience decreases the likelihood of generating a success (extreme right-tail) and increases the likelihood of generating a failure (extreme left-tail) outcome.*

Hypothesis 7 suggests that, within the process of exploration for new knowledge, prior failure experience can simply become a “confusing experience” (Levinthal and March 1993), as innovators blindly explore the vast and rugged *NK* landscape. Recall that exploration to discover superior peaks in the performance landscape corresponds to exploring various technology recombinations across a large number of disparate and interdependent domains (Fleming 2001). High levels of component interdependency suggest that improving the ability to effectively explore new technology recombinations reduces an understanding of the performance impacts of simultaneously modifying these technology configurations, that is, understanding the “complex interdependent co-variations among events” (Reber 1989, Narayanan et al. 2009). By accumulating prior success experience, an innovator gains a better understanding of these interdependencies and more readily identify superior peaks on the rugged performance landscape, with less error, in the future. Consistent with the above arguments, in Hypothesis 9 below, we propose that prior success experience and exploration jointly enhance innovation performance.

***HYPOTHESIS 9:*** *The interaction between exploration and prior success experience increases the likelihood of generating a success (extreme right-tail) and decreases the likelihood of generating a failure (extreme left-tail) outcome.*

### **2.3 Data and Empirical Setting**

The data used for this study is based on the biomedical device industry. Focusing on a single industry allows us to control for cross-industry variations which can account for differences in innovation performance (Klevorick et al. 1995). In addition, the biomedical device industry provides an appropriate context for our study as knowledge creation activities have the potential to generate a wide distribution of innovation outcomes, including innovation breakthroughs, in this industry (Denend and Zenios 2006). In addition, firms in this industry draw upon their prior innovation experience. Our level of analysis is the patent, based on data drawn from the National Bureau of Economic Research (NBER) database (Hall et al. 2001). The data sample includes all patents applied for between 1985 and 1994 within Sub-category 32, which represents innovations related to surgical and medical instruments (Hall et al. 2001). This time window provides five years of data for the focal patents of the study (1990-1994), as well as five years of history for generating our measures to capture prior success and failure experience (1985-1989). Overall, our sample consists of 13,464 focal patents granted to 3390 patent assignees during the specified time period. In addition, for a sub-sample of patents granted to publicly listed firms, we retrieved data from Compustat for firm R&D Intensity and Total Assets, as control variables.

### 2.3.1 Dependent Variables

#### *Innovation Outcome Measures*

The perceived value of an innovation has been shown to be manifested in the number of forward citations received (Albert et al. 1991). Patents which receive a substantial number of forward citations are shown to be associated with the generation of technological breakthroughs, the subsequent commercialization of radical innovations, and the realization of supernormal profits (Trajtenberg 1990, Sorescu et al. 2003). On the other hand, patents which receive very few forward citations are typically reflective of incremental innovations, below-average scientific merit and competitive failures (Tushman and Anderson 1986, Sood and Tellis 2009). In this study, we examine the factors which contribute to success and failure in the knowledge creation process.

For each variable the subscript  $i$  denotes the patent,  $j$  denotes the assignee to which patent  $i$  was granted, and  $t$  denotes the year in which a patent was applied. For each patent, we define the variable  $Success_{ijt}$  as an indicator of whether a particular patent  $i$ , filed by assignee  $j$ , and applied for in year  $t$  is considered a breakthrough innovation, which is evaluated as having above-average technological value (i.e., an extreme right-tail realization), relative to the other innovations in the sample. Specifically, the success variable is a binary variable which equals 1 if patent  $i$  falls in the top 3 % relative to the number of forward citations received by other patents granted in year  $t$  (Ahuja and Lampert 2001, Fleming 2001, Singh and Fleming 2010, Azoulay et al. 2011). We chose the top 3 % as the cutoff for defining an extreme right tail outcome based on a visual inspection of a natural breakpoint in the distribution of the number of forward citations received. Other cutoff points used in previous studies have ranged from the top 1%

(Ahuja and Lampert 2001) to the top 5% (Singh and Fleming 2010). In addition to using the top 3%, for robustness, we also ran the analysis using the top 1% and the top 2 % as alternate cutoff points. The count of forward citations received for the focal patents (1990-1994) are based on the population of patents granted between 1991 and 1999 (Hall et al. 2001). The percentile rankings are done on an annual basis, versus an aggregate basis for the entire sample to account for exposure effects, which lead to earlier patents having more time to accumulate forward citations. For each patent, we define the variable  $Failure_{ijt}$  as an indicator of whether a particular patent  $i$ , filed by assignee  $j$ , and applied for in year  $t$  is evaluated as having below-average technological value (i.e., an extreme left-tail realization), relative to the other innovations in the sample. Specifically, we define failure as a binary variable that equals 1 if patent  $i$  receives zero forward citations. (Singh and Fleming 2010, Azoulay et al. 2011).

### **2.3.2 Independent Variables**

#### ***Exploration and Exploitation***

Consistent with several previous studies on innovation, we consider the details of the patent's backward citations as a proxy for revealing an innovator's knowledge creation activities (Fleming 2001, Ahuja and Lambert 2001, Katila and Ahuja 2002, Nerkar 2003). In order to test the ambidexterity hypothesis, we characterize exploration and exploitation as orthogonal measures (Gupta et al. 2006, Schilling and Green 2011). Currently, no clear consensus exists on a well-accepted definition of exploration versus exploitation (Gupta et al. 2006). However, several proxies exist for distinguishing between exploration to discover a new technological knowledge versus exploitation and reuse of existing knowledge.

We define the variable  $Explore_{ijt}$  as the measure of exploration undertaken in developing patent  $i$ , filed by assignee  $j$ , and applied for in year  $t$ . Based on the  $NK$  model of search and innovation, we define exploration as increasing in the breadth, range and diversity of inputs to the knowledge creation process. This enables the recombination of multiple components (high  $N$ ), across distant or unrelated domains (high  $K$ ) in order to explore and discover new technological trajectories (Fleming and Sorenson 2001, Quintana-Garcia and Benavides-Velasco 2008, Leiponen and Helfat 2009, Singh and Fleming 2010, Gruber et al. 2012, Schilling and Green 2011). Similar to Quintana-Garcia and Benavides-Velasco (2008) we use a Herfindahl diversification index as a proxy for an innovator's degree of exploration (Quintana-Garcia and Benavides-Velasco 2008, Gruber et al. 2012). The diversification index, which ranges from 0 to 1, considers both the breadth and concentration of technology classes referenced in the backward citations for a particular patent. An explore measure that is high (closer to one) suggests a higher degree of search breadth and greater emphasis on accessing a wide range of technology classes when generating the innovation. Conversely, if the patent is based on a very narrow range of technology classes then the measure of exploration is low, or closer to zero.

While we consider exploration as the search for new knowledge, building on Fleming and Sorenson's (2004) notion of "science as map", we define exploitation as the degree to which the innovation search process relies on the reuse and refinement of extant technological knowledge. We define the variable  $Exploit_{ijt}$  as the degree of exploitation undertaken for a particular patent  $i$ , filed by assignee  $j$ , and applied for in year  $t$ . To capture the degree to which an innovation exploits extant knowledge, we consider the

average age of all the contributing technologies utilized for a particular patent as a proxy for the technological maturity of the knowledge base upon which the patent was built (Hamilton and Singh 1992, Ahuja and Lambert 2001, Sorenson and Stuart 2000). Specifically, we measure the average time difference in years, or backward lag, between the application date of the focal citing patent  $i$  and the application date of all the patents which were backward cited as prior art. An exploit measure that is large, this indicates the focal patent placed greater emphasis on exploitation of more mature technologies or was increasingly based on extant technology. The average and maximum backward lag observed in the sample was 11 years and 97 years, respectively.

### ***Prior Innovation Experience***

We consider how the patent assignee's innovation experience contributes to the assignee's future innovation performance. We define the variable *Prior Success*  $_{jt-5}$  to capture the assignee's cumulative success experience. This variable embodies the cumulative number of successful patents, with respect to the number of forward citations received, which were granted to assignee  $j$  over the five year period leading up to year  $t$  of the focal patent. Similarly, we define the variable *Prior Failure*  $_{jt-5}$  as the cumulative number of failed patents, with respect to the number of forward citations received, which were granted to assignee  $j$  over the five year period leading up to year  $t$  of the focal patent. We generate four additional variables to examine the interaction of prior success experience and prior failure experience with exploration and exploitation.

## **2.4 Methods**

We estimate the impact of exploration, exploitation, prior success experience and prior failure experience on the likelihood that a subsequent patent will be a success or a



failure. Since we define success and failure as dichotomous variables, we use a logistic regression model to estimate the effect of the independent variables on the likelihood of success and failure (Singh and Fleming 2010). Our data for estimating the probability of success and failure have a high proportion of zeros and ones. Therefore, we conduct our analysis using the complementary log-log model. To account for temporal effects within the industry, we control for the year in which the patent was applied. We also controlled for the unique patent technology classes (i.e. 128, 600, 601, 602, 604, 606, 607), within Sub-category 32, to which a patent belongs (Hall et al. 2001). In addition, we perform a clustered analysis to account for assignee (firm) level factors which could impact the number of forward citations received by a patent. We applied a logarithmic transformation to adjust for the skewed distribution of the Exploit, Prior Success and Prior Failure variables. To control for multi-collinearity, we mean-centered the exploration, exploitation, prior success and prior failure experience variables and calculated the variance inflation factors across all variables (Aiken and West 1991). For robustness, we ran a zero inflated binomial regression to account for the possibility that a focus on either exploration or exploitation may account for excessive instances of zero forward citations being received by a patent.

We recognize the possibility for endogeneity with respect to factors which may influences both an innovator's choice to explore or to exploit, as well as the number of forward citations received. In our model of knowledge creation (Figure 2.1), we assume that exploration, exploitation, prior success experience and prior failure experience predict the likelihood that an innovation will be a success or a failure. In addition, we consider that prior success and prior failure experience may also affect the degree of

exploration and exploitation undertaken. Given the limitations of the data, no appropriate instrumental variables were available to control for these factors. Therefore, for robustness, we test a seemingly unrelated regression (SUR) mediator/moderator model in which prior success and prior failure experience predicts the degree of exploration and exploitation undertaken, and also moderates the effect of exploration and exploitation on the innovation performance (Zellner 1962, Baron and Kenney 1986).

## **2.5 Results**

Table A1 provides a breakdown of the range of knowledge creation strategies demonstrated across the sample, based on the degree of exploration and exploitation undertaken for each patent. We classify the degree of exploration and exploitation employed for each patent as high or low, based on the level of exploration and exploitation for each patent relative to the median level of exploration and exploitation across the entire sample of patents. From Table A1, we find that 23% of the patents were based on a high exploration-high exploitation strategy, reflecting an ambidextrous knowledge creation strategy. In addition, 15% of the patents reflected a high exploration-low exploitation knowledge creation strategy, reflecting a highly exploratory focus. On the other hand, 33% of the patents exhibited a low exploration-high exploitation knowledge creation strategy, reflecting a highly exploitative focus. Therefore, our sample provides a wide cross-section of knowledge creation strategies against which to test our hypotheses.

Table A2 presents a summary of the descriptive statistics and pair wise correlations of the dependent and independent variables. For the full sample, the minimum, average and maximum of the exploration variable is 0, 0.38 and 0.91,

respectively. The minimum, average and maximum of the exploitation variable is 0, 11.02 and 97 years, respectively. Approximately 3% of the patents were classified as a success, reflecting an extreme right tail outcome, whereas approximately 10% of the patents were classified as a failure, reflecting an extreme left tail outcome. For those patents ranked as a success, we examined the proportion of self-citations as well as the proportion of citations received from a category external to the major patent category 3 assigned for drugs and medical (Hall et al. 2001). We found that for the success patents, as compared to the full sample of patents, the average percentage of self-citations were comparable at approximately 10%. For the full sample approximately 14% of the forward citations received came from an external patent category, whereas for the success patents approximately 7% of the forward citations received came from an external patent category. The minimum, average and maximum number of prior successes experienced by an assignee, in the five years preceding the focal patent, is 0, 1.52 and 22, respectively. While, the minimum, average and maximum number of prior failures experienced by an assignee, in the five years preceding the focal patent, is 0, 1.01 and 17, respectively.

Table A3 provides the results of the complementary log-log model which estimates the effects of exploration, exploitation, prior success experience and prior failure experience on the likelihood of success and failure. Models S1, S2, S3, S4 and S5 test the impact of the independent variables on the likelihood of success. Similarly, models F1, F2, F3, F4 and F5 test the impact of the independent variables on the likelihood of failure. Therefore, for models S1-S5, a positive coefficient indicates an increased probability of success and hence a performance benefit. Conversely, for models

F1-F5, a negative coefficient indicates a reduced probability of failure and hence a performance benefit.

From Table A3, models S1 and F1 includes the exploration and exploitation variables. Recall that Hypothesis 1 conjectured that exploration is positively related to an increased likelihood of both successes and failures. The positive coefficient for exploration in model S1 indicates that exploration increases the likelihood of success ( $\beta = 0.34$ ,  $p < 0.001$ ). Whereas, the negative coefficient for exploration in model F1 indicates that exploration reduces the likelihood of failure ( $\beta = -0.04$ ,  $p < 0.10$ ). Hypothesis 2 posited that exploitation is related to a decreased likelihood of both successes and failures. The negative coefficient for exploitation in model S1 indicates that exploitation decreases the likelihood of success ( $\beta = -0.49$ ,  $p < 0.001$ ). Whereas, the positive coefficient for exploitation in model F1 indicates that exploitation increases the likelihood of failure ( $\beta = 0.33$ ,  $p < 0.001$ ). Therefore, Hypotheses 1 and 2 are partially supported, with respect to the extreme right tail outcomes (success), but not with respect to the extreme left tail outcomes (failure).

Our findings deviate from prior conceptual assumptions about the impact of exploration and exploitation on the probability of generating extreme left-tail outcomes from a distribution of innovation outcomes. Exploration and exploitation have been defined as having variance-increasing and variance-reducing impacts on innovation performance outcomes, respectively (Fleming 2001, He and Wong 2004). Furthermore, it is typically assumed that exploration (exploitation) symmetrically increases (decreases) the likelihood of both left-tail and right-tail outcomes (March 1991, Fleming 2001, Gupta et al. 2006). However, contrary to the expected results, our findings indicate that

exploration increases the likelihood of success, but also decreases the likelihood of failures. While we show that the direct effect of exploitation is to reduce the likelihood of success, but to increase the likelihood of failure. Therefore, we conduct three additional tests. In the first two additional analyses, we normalize the number of forward citations received by a patent, relative to all other patents granted in the same year, to have a mean of zero and a standard deviation of one. We then run a linear regression to estimate the impact of all the variables in the study on both the normalized level of citations as well as the absolute deviation. This analysis allows us to examine the impact of the independent variables on changes to both the mean as well as the variance of the distribution of the number of forward citations received (Taylor and Greve 2006). In a third test, we conduct a quantile regression analysis on the number of forward citations received (Singh and Fleming 2010). Unlike linear regressions, which estimate the impact of an independent variable on the mean, quantile regressions estimate the effect of an independent variable on the magnitude of each percentile of a distribution (Singh and Fleming 2010).

The results of the linear regression on the mean and variance of the number of forward citations received are presented in Table A4. The quantile regression results are provided in Table A5. From model V1 of Table A4 we see that exploration increases the variance in the number of citations received ( $\beta = 0.40$ ,  $p < 0.001$ ). While from model M1 of Table A4, we find exploration also increases the mean ( $\beta = 0.06$ ,  $p < 0.001$ ). On the other hand, from model V1 of Table A4 we see that exploitation decreases the variance ( $\beta = -0.05$ ,  $p < 0.001$ ). While model M1 of Table A5 shows that exploitation also decreases the mean ( $\beta = -0.14$ ,  $p < 0.001$ ). These additional mean/variance regression analyses provide important insights. Although exploration increases the variance and

uncertainty of outcomes, as expected, it also increases the mean (a right shift of the distribution), so that the likelihood of an extreme left-tail outcome, that is the risk of failure, is reduced. Conversely, although exploitation decreases the variance and the uncertainty of outcomes, given that the mean also decreases (a left shift of the distribution) then the likelihood of an extreme left-tail outcome (the risk of failure) is amplified. The quantile regression coefficients for the lower (10<sup>th</sup>) , middle (50<sup>th</sup>) and upper (90<sup>th</sup>) percentiles are given in Table A5. The direction and relative magnitude of the quantile coefficients confirms the mean/variance findings.

Although we find that exploitation neither directly improves success nor reduces failure, we also consider that exploitation may be undertaken to benefit from the synergies related to ambidexterity, as proposed in Hypothesis 3. We include the interaction term between exploration and exploitation in models S2 and F2 of Table A3. The results from model F2 indicate that jointly undertaking exploration and exploitation increases the likelihood of failure ( $\beta = 0.11$ ,  $p < 0.001$ ). Therefore, our results provide support for the alternate Hypothesis 3B which highlights the challenges of ambidexterity. This result provides an indication of the potential coordination costs and negative performance implications which may result when trying to combine these two incompatible modes of knowledge creation (March 1991, Leonard-Barton 1992, Gilson et al. 2005, Raisch and Birkinshaw 2008). The interaction effect for the likelihood of success is not significant.

From Table A3, models S3 and F3 includes only the prior success and prior failure variables. In Hypothesis 4, we propose that learning from success may improve performance. In model S3, the positive coefficient for the effect of prior success

experience on the likelihood of subsequent success indicates that innovators who have succeeded in the past are more likely to generate a breakthrough in the future ( $\beta = 0.59$ ,  $p < 0.001$ ). Furthermore, from model F3, prior success experience provides a dual benefit, as it also reduces the likelihood of future failures ( $\beta = -0.19$ ,  $p < 0.001$ ). This suggests that, as they search over the rugged landscape, innovators who have successfully located a peak on the *NK* landscape in the past are more likely to successfully generate a breakthrough in the future, with less error (Sorenson et al. 2006).

In Hypothesis 5A, we propose that the opportunity to learn from prior failures may improve performance. However, from model S3, we find that prior failure experience reduces the likelihood of future success ( $\beta = -0.54$ ,  $p < 0.001$ ). Furthermore, from model F3, we find that prior failure experience increases the likelihood of subsequent failures ( $\beta = 0.10$ ,  $p < 0.01$ ). These findings do not support a theory of learning from failure, as posited in Hypothesis 5A. To the contrary, we find support for the alternate hypothesis, given in 5B, that prior failure experience degrades future performance. This result suggests that innovators may become trapped in a valley during search on the rugged *NK* landscape (Sorenson et al. 2006). Collectively, the results of Hypotheses 4 and 5 are also consistent with the suggestion, based on attribution theory, that it may be easier for innovators to learn and improve performance based on their prior success experience versus their failure experience (KC et al. 2013).

We probe further into the benefits of prior experience to consider how this knowledge moderates the effect of future exploration and exploitation. We include the interaction terms between exploration, exploitation, prior success and prior failure in models in S4 and F4 (Table A3). However, we find that in several cases the sign of the

coefficient for the interaction term changes in the full models, given in columns S5 and F5 of Table A3. For robustness, given the possibility for endogeneity, we test a seemingly unrelated regression (SUR) mediator/moderator model. In this model we assume that prior success and prior failure experience predict the likelihood of future success and failure, as well as moderate, not only, the degree, but also the effectiveness, of exploration and exploitation. We find that the results of the interaction effects between prior innovation experience and exploration, as well as exploitation from models S4 and F4 (Table A3) are consistent with the results of model 3 (SUR Mean, Table A6) and model 6 (SUR Variance, Table A6), which accounts for this endogenous relationship. Therefore, we focus on the interaction results of models S4 and F4, versus models S5 and F5 from Table A3.

The results of model S4 (Table A3) support Hypothesis 6 and demonstrates that the interaction between exploitation and prior failures increases the likelihood of a breakthrough success ( $\beta = 0.24$ ,  $p < 0.001$ ). This is an interesting finding as the earlier results of Hypotheses 2 and 5B show that the direct effects of both exploitation and prior failure experience is to decrease the likelihood of success. However, the results of Hypothesis 6 show that an innovator is able to reverse these effects and is able to derive a positive indirect learning benefit from jointly leveraging both exploitation and prior failure experience. Collectively, the results of Hypotheses 2, 4 and 6 demonstrate that prior failure experience is necessary, but not sufficient for learning from failure to occur. Specifically, the results suggests that it is also necessary for an innovator to leverage the mature, well-codified knowledge, associated with exploitation (Narayanan et al. 2009), in order to extract the information embedded in prior failure and thereby improve future



performance. On the other hand, we also find support for a potential “exploration failure trap” (Levinthal and March 1993), as posited in Hypothesis 7. Although failure experience, in conjunction with exploitation, improves performance, we find that the interaction between exploration and prior failures decreases the likelihood of a breakthrough success ( $\beta = -0.16$ ,  $p < 0.10$ ). Consequently, prior failure experience is a potential “two-edged sword”, depending on whether the future search activities are pursued with an exploratory or exploitative focus. Although innovators pursuing search with an exploitative focus may benefit from learning from failure, failure experience may also lead to confusion when search activities have a more exploratory focus.

We also find partial support for an “exploitation success trap” as posited in Hypothesis 8. From Table A3, model S1 and Table A4, model V5 we found that the main effect of exploitation is to reduce the likelihood of breakthrough success and reduce uncertainty, respectively. Additionally, from model S4 (Table A3), we find that the interaction between exploitation and prior success further decreases the likelihood of a breakthrough success ( $\beta = -0.23$ ,  $p < 0.01$ ). Furthermore, from model V4 of Table A4, we find that prior success experience makes exploitation more efficient at reducing the variance of innovation outcomes ( $\beta = -0.04$ ,  $p < 0.05$ ). Our results, therefore, provide empirical support for the notion of a success learning trap (Levinthal and March 1993), as prior success experience reinforces exploitation’s main effect of generating predictable, but sub-optimal, performance outcomes (Levinthal and March 1993, Leonard-Barton 1992, Audia and Goncalo 2007). Based on models S4 and F4 (Table A3), the interaction effect between prior success experience and exploration, as proposed in Hypothesis 9 is not supported.

### **2.5.1 Robustness Tests**

#### ***Alternate Success Percentile Cut-off Points***

For robustness, we ran the analysis for the likelihood of breakthrough success using three alternate success percentile cut-off points. The results of the complementary log-log model comparing the likelihood of being in the Top 1%, 2% and 3%, which are given in Table A7, are consistent with the main results from Table A3.

#### ***Zero Inflated Negative Binomial Model***

Since the number of forward citations received is truncated at zero, we ran the analysis for the mean number of citations received, using a zero inflated negative binomial model. Even accounting for the inflation, the results of the zero inflated negative binomial model, given in Table A8, are consistent with the main results.

#### ***Sub-sample of Publicly Listed Firms***

For a sub-sample of patents granted to publicly listed firms, we matched patent assignees to firms, in order to retrieve financial data from Compustat (Bessen 2009). We ran the analysis for the sub-sample, including controls for the magnitude of the firm's assets as well as for R&D intensity. The results for the subsample are given in Table A9. The subsample includes 2835 focal patents assigned across 125 publicly listed firms. As expected asset size and R&D intensity were positively related to the likelihood of breakthroughs, but had a non-significant impact on the likelihood of failure. For the sub-sample the results for the impact of exploration, exploitation, prior failure experience and prior success experience were all consistent with the main results on the likelihood of breakthrough success and failure, although the impact of exploration was non-significant. Interestingly, we do not find support for learning from failure, either with respect to the

main effect of prior failure experience, or through the interaction terms. This suggests that more mature publicly-listed firms may lack either the motivation or ability to leverage their failure experience.

### ***Alternate Exploitation Measure***

We recognize that no clear consensus exists on a well-accepted definition of exploration versus exploitation (Gupta et al. 2006). In selecting the measures for our study we consider two key features of the knowledge creation process which we wish to examine: (i) the impact of knowledge creation activities on the uncertainty of innovation outcomes and (ii) the complementary relationship between exploration, exploitation and prior innovation experience. Specifically, we consider exploration as experimentation with diverse components to discover a new technological trajectory versus exploitation as the reuse of components from an existing technological trajectory (Dosi 1982, Fleming 2001, Linton and Walsh 2004). Our notion of exploitation is consistent with Benner and Tushman's definition of exploitation as search activities which "involve improvements in existing components and build on the existing technological trajectory" (2002, pg. 679)

This definition of exploitation as building on an existing technological trajectory (Benner and Tushman 2000), or relying on extant science as a "map" of the technological landscape (Fleming and Sorenson 2004), differs from some studies in which exploitation has been defined as the depth of an innovator's existing knowledge base. Studies invoking this definition have measured the degree of exploitation as: "accumulation of search experience with the same knowledge elements" (Katila and Ahuja 2002); "knowledge stock", "absorptive capacity" (Wu and Shanley 2009); and "ongoing use of a firm's knowledge base" (Vermeulen and Barkema 2001). Therefore, for robustness we

defined an alternate exploitation variable, *Exploit II*, to capture the depth of an innovator's existing knowledge stock, with respect to the knowledge elements used in the focal patent. Similar to Fleming (2001), we calculated the number of times, within the prior ten years, that the assignee utilized the same configuration of technological classes in prior patents. This captures the assignee's cumulative familiarity with this technology combination. The results for the likelihood of success and failure models using this alternate exploitation measure are given in Tables A10 and A11.

In our model of induced and autonomous learning (Figure 2.1) we explicitly disaggregate the two elements of exploitation referenced above by considering both: (i) induced learning based on exploitative search, which builds on an existing technological trajectory, and (ii) autonomous learning which leverages the knowledge stock gained from cumulative search experience. Given that we already consider an innovator's accumulated success experience and failure experience, additionally considering exploitation as the reuse of the innovator's knowledge stock is redundant. The results from Tables A10 and A11 confirm this. Table A10 shows that the main effect of prior failure experience is to reduce the likelihood of success, and that the interaction between prior failure experience and cumulative experience with these combinations (*Exploit II*) further reduces this likelihood of success. The interaction between prior failure experience and cumulative combination experience (*Exploit II*) suggests that experience with the same technology combination does not improve performance, given prior failure experience. Analogously, Table A10 shows that the main effect of success experience is to increase the likelihood of success and that the interaction between prior success experience and cumulative experience with these combinations (*Exploit II*)

further increases this likelihood of success. Since we hypothesize an interaction between exploitation and prior innovation experience, an exploitation measure which also reflects cumulative search experience is therefore a redundant measure.

### ***Discounting Prior Experience***

Given the importance of prior experience as both an independent variable and a moderator in the model, for robustness we ran the analysis using two alternate specifications for discounting prior success experience and prior failure experience. First, in Table A12 we discounted both prior success experience and prior failure experience by the square root of the age of the experience. While, in Tables A13 we implemented a linear discounting factor (Baum and Ingram 1998, Haunschild and Sullivan 2002). The results after discounting the effects of prior experience are consistent with the main results from Table A3, although there are a few changes in the level of significance. For example, comparing Table A3 with Tables A12 and A13, the interaction term between prior failure experience and exploitation, still positively impacts the likelihood of success, however, the significance is slightly reduced in the discounted models. In one case, a term which was significant in Table A3 is no longer significant in Tables A12 and A13, and that occurs with the estimation of the impact of prior failure experience on increasing the likelihood of a subsequent failure

## **2.6 Implications for Theory and Practice**

Our results address important open questions in the research on ambidexterity, with respect to understanding , “when ambidexterity is more or less useful?” and “when do the benefits of ambidexterity outweigh the costs?” (O’Reilly and Tushman 2013, p.26). Specifically, we contribute to a more nuanced understanding of both the short-

term and long-term challenges, and benefits, of pursuing an ambidextrous knowledge creation strategy during the innovation process. We provide some important findings for advancing a more comprehensive model of how innovators benefit from exploration and exploitation, as well as how they can learn from prior failure and success experience, in order to increase the likelihood of generating a breakthrough innovation and mitigate the risks of failures. With these insights, managers can become better informed regarding strategies for managing two important categories of “risk” in the innovation process: the risk of “sinking the boat”, that is, the risk of failure, versus the risk of “missing the boat”, which is the risk of missing the opportunity for successfully generating a breakthrough (Dickson and Giglierano 1986).

Our results have important theoretical and managerial implications. Firstly, we highlight the importance of making a critical conceptual distinction between classifying exploration as an uncertain versus a risky undertaking. Our findings weaken support for the traditional classification of exploration as a risk-increasing activity, since exploration increase the upside potential for generating breakthrough successes, while decreasing the downside risk of failures. Traditionally, it has been thought that exploration leads to more uncertainty, as well as higher probability of failure. However we find that while exploration does increase the variance of innovation outcomes, this occurs through an asymmetric upward shift in the distribution (Singh and Fleming 2010). Specifically, we find that exploration increases both the mean as well as the variance of innovation outcomes. This is consistent with March’s (1991) view on the performance implications of the exploration of new knowledge: *“Some learning processes increase both average performance and variability. A standard example would be the short-run consequences*

*from adoption of a new technology. If a new technology is so clearly superior as to overcome the disadvantages of unfamiliarity with it, it will offer a higher expected value than the old technology. At the same time, the limited experience with the new technology (relative to experience with the old) will lead to an increased variance”(p. 83).*

In addition, while exploitation has been touted for its risk-reducing benefits, we do not find significant support for this claim. While we do find that exploitation reduces the variance of innovation outcomes; we find that this occurs through an asymmetric downward shift in the distribution (Singh and Fleming 2010). Therefore, while the absolute uncertainty of outcomes is reduced with exploitation, the likelihood of generating breakthrough outcomes is reduced, and the likelihood of generating breakthrough failures is increased. If an organization’s objective is to reduce uncertainty, then exploitation does provide uncertainty-reduction benefits. However, if the goal is risk-reduction, that is to improve the upside potential for generating breakthroughs and/or to reduce the downside risk of failures, then exploitation, as defined in this study, does not achieve this objective. Therefore, we demonstrate that exploitation can lead to predictable, but sub-optimal performance (Levinthal and March 1993, Leonard-Barton 1992).

While traditionally studies on learning have focused on repetitive, production-related tasks, our results demonstrate that non-repetitive tasks, such as innovation, can also benefit from learning-by-doing. We find support for the theory of learning from success, as our results demonstrate that prior breakthrough experience leads to a higher likelihood of generating breakthroughs in the future. However, we also find that innovators can succumb to an “exploitation success trap”. As success experience and

exploitation jointly reduce uncertainty, this could lead to an overinvestment in exploitation because “*organizations discover the short-term virtue of local refinement and the folly of exploration. As they develop greater and greater competence at a particular activity, they engage in that activity more, thus further increasing competence and the opportunity cost of exploration*” (p. 106).

While we find consistent evidence of the effectiveness of learning from success, we did not find support for the main effect of learning from failure. Furthermore, we show that innovators are prone to falling into an “exploration failure trap”, as demonstrated by the interaction between exploration and prior failure experience, which jointly reduce the likelihood of future breakthroughs (Levinthal and March 1993). Cannon and Edmondson (2005) suggest that, in order to benefit from prior failure experience, innovators must also have the requisite knowledge and processes in place, to be able to detect and analyze these failures. Consistent with this suggestion, we find that only with a focus on exploitative search, based on mature, well-codified knowledge, can innovators jointly leverage both exploitation and prior failure experience, in order to increase the likelihood of generating a future breakthrough. This result clearly demonstrates the benefit of pursuing ambidexterity. Although exploitation has a negative main effect on the likelihood of generating breakthroughs, innovators still need to balance exploration with exploitation, in order to leverage the knowledge available from prior failure experience. Importantly, therefore, we establish a critical link between the theory of ambidexterity and the theory of learning from failure. Whereas these two theories have been offered as alternative views on dealing with the risks associated with



exploration, our findings demonstrate that, in fact, pursuing exploitation and learning from failure are complementary processes, which operate in tandem with each other.

In spite of demonstrating the benefits of pursuing exploitation, in order to extract the information embedded in past failures, our results also demonstrate the paradoxical challenges associated with simultaneously undertaking both exploration and exploitation. As opposed to those studies which demonstrate the benefits of ambidexterity (Katila and Ahuja 2002, Rothaermel and Deeds 2004, Cao et al. 2009, Chandrasekaran et al. 2011), our results indicate the potential negative performance implications of ambidexterity, as reflected by the negative interaction effect between exploration with exploitation. This result provides empirical support the notion that exploration and exploitation are incompatible knowledge creation activities, which may be difficult to combine (Leonard - Barton 1992, Gupta et al. 2006, Raisch and Birkinshaw 2008, Andriopoulos and Lewis 2009). Zhou (2011) finds that pursuing multiple related activities can provide synergies, however he notes that the cost of coordination may dominate any synergistic benefits. Similarly, our findings suggest that innovators which pursue an ambidextrous strategy must recognize both the potential perils as well as the possible benefits of such a strategy. Figure 1.2 below summarizes both the positive and negative performance implications which result from the main effects as well as the interaction effects between exploration, exploitation, prior success and prior failure experience (Table A3, Models S1-4/F1-4; Table A4, Model V4, Table A6 Models 3 and 6). Figure 1.2 highlights ways in which exploration, exploitation, success experience and failure experience deliver innovation performance benefits which are both conflicting and synergistic. Collectively, our results illustrate that pursuing an ambidextrous knowledge creation strategy necessitates a

delicate balancing act, in order to manage the short-term and long-term benefits, and perils, of exploration, exploitation, learning from success and learning from failure.

<b>VARIANCE OF INNOVATION OUTCOMES</b>	<b>Exploration</b> Increases Variance	<b>Exploitation</b> Decreases Variance
<b>Prior Success Experience</b> Increases Variance	Interaction Not Significant	Decreases Variance <i>(EXPLOITATION SUCCESS TRAP)</i>
<b>Prior Failure Experience</b> Decreases Variance	Decreases Variance	Increases Variance <i>(LEARNING FROM FAILURE ENABLED WITH EXPLOITATION)</i>

<b>LIKELIHOOD OF SUCCESS</b>	<b>Exploration</b> Increases Success	<b>Exploitation</b> Decreases Success
<b>Prior Success Experience</b> Increases Success	Interaction Not Significant	Decreases Success <i>(EXPLOITATION SUCCESS TRAP)</i>
<b>Prior Failure Experience</b> Decreases Success	Decreases Success <i>(EXPLORATION FAILURE TRAP)</i>	Increases Success <i>(LEARNING FROM FAILURE ENABLED WITH EXPLOITATION)</i>

<b>LIKELIHOOD OF FAILURE</b>	<b>Exploration</b> Decreases Failure	<b>Exploitation</b> Increases Failure
<b>Prior Success Experience</b> Decreases Failure	Interaction Not Significant	Interaction Not Significant
<b>Prior Failure Experience</b> Increases Failure	Interaction Not Significant	Interaction Not Significant

**Figure 2.2: Interdependence of Exploration, Exploitation and Prior Experience**

## **2.7 Future Research**

Our current study has focused on exploration and exploitation knowledge creation activities as determinants of the performance outcomes from the innovation process. These findings have managerial implications for firms which must engage in the uncertain process of knowledge creation and innovation. However, we recognize several limitations of this study. First, by focusing on the biomedical device industry, we are unable to necessarily generalize the results to all industries. Second, we study the impact of exploration, exploitation and ambidexterity on success and failure, once a patent has been granted at the invention stage of the innovation process. However, we recognize that the innovation process is multistage process, (Krishnan and Loch 2005, Gaimon and Bailey 2012), for which success and failure outcomes are possible outcomes at each stage of the process (Girotra et al. 2007, Hora and Dutta 2012). Firstly, a patent application could be considered a success or failure, depending on whether or not the patent is successfully approved and granted. Thereafter, conditional on being granted, a patent's success or failure can be measured as a function it's demonstrated technical value and scientific merit. Once product development and commercialization begins, additional measures of innovation success and failure include the number of new products created, product quality, frequency of product recalls, level of market profitability and other innovation-related financial metrics (Griffin and Page 1993, Katila and Ahuja 2002, Sorescu et al. 2003, Artz et al. 2010). Although our results identify the impact of exploration and exploitation on the outcomes in the invention stage, our study does not consider the impact of these knowledge creation activities on the preceding (patent application) and subsequent stages (commercialization) in the innovation process.

Therefore, the results of our study may underestimate or overestimate the impact of exploration, exploitation and ambidexterity on overall success and failure in the innovation process.

The fact that this study considers only a single stage in the multistage innovation process also provides several opportunities for future research. Whereas, in this paper we demonstrate the positive benefits of exploration on generating a scientific breakthrough, a follow-up step would be to examine the impact of exploration, exploitation and ambidexterity on both the preceding and subsequent steps in the innovation process. To our knowledge, no study has examined the impact of exploration and exploitation on the likelihood of success versus failure in the patent application process. We have begun a study to extend our findings on the impact of exploration and exploitation, in the patent application process. Secondly, we are also interested in studying the relationship between exploration, exploitation and ambidexterity on the subsequent commercialization stage, to examine the impact of knowledge creation activities on the eventual technical success and failure of the innovation. In the case of the biomedical industry, technical success and failure would be reflected in the number of medical device recalls. This is especially important in the medical device industry where, although the potential to develop breakthrough innovations is important, the final product quality and the likelihood of technical failures and product recalls are also critical. In low-risk medical device sectors such as orthopedics, a high level of technical uncertainty may be more acceptable (Denend and Zenios 2006). However, in more risk-sensitive segments, for example cardiac related innovations, firms may be less prone to undertake exploration of new knowledge, to generate breakthrough innovations, if this can lead to a higher risk of

technical failure and recalls in the subsequent product development and commercialization stages. Finally, research opportunities also exist for studying the creation and diffusion of knowledge across multiple industries. We find that for the success patents approximately 7% of the forward citations received came from a patent category external to the patent category of the focal patent. This suggests that another fruitful area for research would be to examine the knowledge spillover implications with respect to how exploration in a one industry may provide opportunities for exploration in related and unrelated technological domains (Autio et al. 2004, Yang et al. 2010)

## **CHAPTER 3**

### **A DYNAMIC MODEL OF KNOWLEDGE CREATION FROM EXPLORATION AND EXPLOITATION**

#### **3.1 Introduction**

A critical challenge faced by firms who achieve their competitive advantage through innovation is to identify efficient and effective strategies for managing the high levels of technical uncertainty associated with the process of innovation. Technical uncertainty, as embodied in the variance of an innovation's performance outcome, reflects the extent of the range and unpredictability of an innovation's performance outcomes. Existing research on the innovation process has utilized the stage-gate and innovation tournament concepts to describe the dynamic evolution of uncertainty throughout the innovation process (Cooper 1990, Sethi and Iqbal 2008, Terwiesch and Ulrich 2009). While both the stage-gate and tournament funnel approaches recommend that uncertainty be continuously reduced over time, some scholars point to the potential pitfall of focusing on resolving uncertainty too early in the innovation process. To the contrary, it has been suggested that generating uncertainty and retaining ambiguity during in the innovation process may be beneficial (Leonardi 2011, Austin et al. 2012). Commenting on the optimal timing of uncertainty resolution during innovation process, Cross (2000) notes, "normally, the overall aim of a design strategy will be to converge on a final, evaluated and detailed design proposal, but within the process of reaching that final design there will be times when it will be appropriate and necessary to diverge, to widen the search or to seek new ideas and starting points" (p. 186). Contrary to the stage-

gate view of continuous uncertainty reduction, this comment suggests that it may be desirable to increase the variance of innovation outcomes at certain points during in the innovation process. These two opposing views suggest that the optimal dynamic evolution and resolution of uncertainty over time during the innovation process remains an open issue (Koput 1997, Leonardi 2011, Bingham and Davis 2012).

Several factors have been suggested to influence the optimal rate and timing of the evolution and resolution of uncertainty during the innovation process including the relative costs and information generation efficiencies across various innovation activities (Thomke 1998, Thomke 2000, Fixson and Marion 2012, Austin et al. 2012), an organization's risk preferences (Carrillo and Gaimon 2004, Sommer and Loch 2009, Manso 2011, Chandrasekaran and Mishra 2012), and differences in short-term versus long-term innovation objectives (Van de Ven and Polley 1992, Das and Teng 1997, Fried and Slowik 2004, Manso 2011). For example, Thomke (1998) suggests that changes in the relative efficiencies of innovation technologies, such as the availability of cheaper prototyping methods for exploring new innovation solutions, may make it optimal to “fail early, fail cheap”, resulting in an increased focus on generating uncertainty early in the innovation process (Thomke 2001, Austin et al. 2012). However, reducing the marginal cost of innovation and experimentation technologies may also have the opposite effect. Fixson and Marion (2012) empirically find that the reduced cost of digital design tools can result in postponing exploratory design iterations, which naturally results in the generation of uncertainty later in the innovation process. In this paper, we introduce a dynamic model which examines the impact of the above-mentioned drivers on the optimal evolution and resolution of technical uncertainty during the innovation process.

To examine the optimal dynamic evolution and resolution of technical uncertainty throughout the innovation process we introduce a dynamic model of knowledge creation. We build on our findings presented in Chapter 2, in which we empirically demonstrate the distinction between exploration and exploitation, as variance-generating versus variance-reducing knowledge creation activities (March 1991, Fleming 2001, He and Wong 2004). March (1991) suggests that pursuing either exploration or exploitation independently, or excessively, can lead to negative performance outcomes. Specifically, as demonstrated in Chapter 2, excessive exploration can result in an innovation which is novel, but with a high degree of uncertainty, whereas excessive exploitation may yield predictable outcomes, but with only marginal performance advantages over the current solutions (March 1991). Avoiding either of these extreme scenarios requires that firms follow an appropriate ambidexterity strategy, which refers to the simultaneous pursuit of both exploration and exploitation.

Several strategies have been proposed for dealing with the challenge of investing in exploration and exploitation simultaneously (Raisch and Birkinshaw 2008, Lavie et. al 2010, O'Reilly and Tushman 2013). One such approach is a *temporal ambidexterity* strategy in which an organization balances exploration and exploitation by sequentially focusing on generating versus resolving uncertainty, at different points in time during the innovation process (Brown and Eisenhardt 1997, Siggelkow and Levinthal 2003). However, the specific optimal sequence of exploration and exploitation activities which should be adhered to when pursuing a temporal ambidexterity strategy remains an open issue (Raisch and Birkinshaw 2008, Raisch et al. 2009, Lavie et al. 2010).



This paper contributes to the theory of temporal ambidexterity by introducing an analytical model which examines the optimal timing of exploration and exploitation knowledge creation activities throughout the innovation process. Our research provides three key contributions to the academic literature, as well as to managers seeking a better understanding of effective temporal ambidexterity strategies (Brown and Eisenhardt 1997, Siggelkow and Levinthal 2003, Andriopoulos and Lewis 2009, Lavie et al. 2010, Smith et al. 2010). Our first contribution is to extend March's (1991) conceptual framework by introducing an analytical model which captures the interdependencies and tradeoffs that arise when simultaneously investing in both exploration and exploitation during the innovation process. Secondly, the existing literature highlights the path-dependent process of knowledge creation. Based on this notion, prior knowledge creation activities, in other words cumulative experience, may impact the effectiveness and hence the optimal sequencing of the knowledge creation activities which follow (Shane 2000, Audia and Goncalo 2007, Eesley and Roberts 2010, Ozkan et al. 2009, Gaimon and Bailey 2012, Argote 2013). Therefore, a key feature of our model is the inclusion of the effect of absorptive capacity, which we characterize as reflecting the impact of prior knowledge creation activities and cumulative experience, on the efficiency and effectiveness of future exploration and exploitation activities (Cohen and Levinthal 1990). Importantly, we extend the generalized model of absorptive capacity and we distinguish between the effects of cumulative exploration versus cumulative exploitation experience, which have differing effects on the effectiveness of future knowledge creation activities. Third, we consider how the optimal investments in knowledge creation are impacted by the decision-maker's short-term versus long-term risk

preferences (Das and Teng 1997, Fried and Slowik 2004), which are functions of the levels of technical uncertainty which remain unresolved at various stages of the innovation process (Van de Ven and Polley 1992, Lenfle and Loch 2010, Azuolay et al. 2011). Incorporating dynamic risk preferences in our model of exploration and exploitation is critical, and leads to new and important insights on the inter-temporal tradeoffs which may occur when simultaneously pursuing exploration and exploitation.

### **3.2 Related Literature**

#### **3.2.1 Exploration, Exploitation and Temporal Ambidexterity**

It is widely accepted that superior innovation performance is realized by firms which can successfully engage in an ambidextrous approach, simultaneously employing both exploration and exploitation. To date the existing literature on ambidexterity has focused on the development of conceptual frameworks (Gupta et al. 2006, Lavie et al. 2010), simulations (Levinthal and March 1981, March 1991), and empirical studies which examine the benefits of employing an ambidextrous knowledge creation strategy during the innovation process (Katila and Ahuja 2002, Rothaermel and Deeds 2004, He and Wong 2004). For example, Katila and Ahuja (2002) empirically find a significant positive interaction effect between exploration and exploitation on the total number of new products introduced by firms pursuing innovation.

Several strategies have been proposed to simultaneously undertake both of these knowledge creation activities. The two strategies most prominently discussed in the literature are (i) organizational (functional/domain) ambidexterity and (ii) temporal ambidexterity (Brown and Eisenhardt 1997, Siggelkow and Levinthal 2003, Raisch and Birkinshaw 2008, Lavie et al. 2010, O'Reilly and Tushman 2013). Firms which

implement an organizational or functional ambidexterity strategy have separate organizational units which focus exclusively on exploration versus exploitation activities (Raisch and Birkinshaw 2008, Lavie et al. 2010). Alternatively, firms which implement a temporal ambidexterity strategy pursue exploration and exploitation simultaneously by sequentially focusing on each mode of knowledge creation at different points in time, within a single organizational unit (Lavie et al. 2010). However, the literature remains unclear as to what specific sequence of activities is most effective when executing a temporal ambidexterity strategy (Bingham and Davis 2012). In one notable exception to this gap in the literature, Rothaermel and Deeds (2004) empirically finds that in the biotech industry, a sequential strategy of exploration followed by exploitation is optimal in order to improve innovation performance, as measured by the total number of new products generated by a firm. In contrast to Rothaermel and Deeds' empirical research study, in our paper we propose an analytical model of temporal ambidexterity, which allows us to characterize the conditions under which the typical explore-exploit as well as the atypical exploit-explore optimal dynamic knowledge creation strategies may optimally arise.

### **3.2.2 Generating and Resolving Uncertainty in Innovation**

Based on empirical research, the importance of considering the benefits of simultaneously pursuing exploration and exploitation is well accepted (Katila and Ahuja 2002, Rothaermel and Deeds 2004, He and Wong 2004, Raisch and Birkinshaw 2008, Lavie et al. 2010). However, the existing normative literature has largely treated the problem of the optimal investment in exploration versus exploitation independently. That is, one subset of the normative literature has considered models of investment in the

variance-generating process of exploration related to search, concept generation and innovation discovery (Gavetti and Levinthal 2000, Huchzermeier and Loch 2001, Kavadias and Sommer 2009, Erat and Krishnan 2012) whereas another subset of the normative literature has considered models of investment in the variance-reducing process of exploitation related to concept selection, testing and uncertainty resolution (Krishnan et al. 1997, Loch and Terwiesch 1998, Loch et al. 2001, Thomke and Bell 2001, Erat and Kavadias 2008). Moreover, the latter body of research, related to concept selection and uncertainty resolution, generally characterizes technical uncertainty as exogenous and learning as a strictly uncertainty reducing strategy. Our paper contributes to both these streams of literature. However, our work differs from the existing literature in that we introduce a model which simultaneously considers the dynamic generation, as well as the dynamic resolution of uncertainty in the innovation process, both of which are endogenous. We model exploration-based learning as generating uncertainty in the innovation process, while exploitation-based learning is associated with uncertainty resolution.

### **3.2.3 Dynamic Decision Making Under Uncertainty**

While the analytical research noted above which mainly considers uncertainty as exogenous, the empirical literature on risk-taking reflects the endogenous nature of uncertainty. However, the empirical literature provides inconsistent findings with regards to when it is optimal to undertake additional risk and to increase uncertainty (Sitkin and Pablo 1992, Forlani and Mullins 2000). Some evidence suggests that decision-makers undertake more risky decisions, such as exploration, to avoid poor performance (March and Shapira 1992, Wiseman and Bromiley 1996). In contrast, other research suggests that

decision-makers undertake less risky decisions, such as investing in exploitation, to preserve deteriorating performance (Staw et al. 1981, March and Shapira 1992). To examine this open issue, regarding optimal risk-taking, we consider how a manager's performance objectives drive the decision to invest in variance-generating exploration versus variance-reducing exploitation. Importantly, we recognize that a manager may have dynamic risk-preferences with respect to his short-term versus his long-term objectives (Das and Teng 1997, Fried and Slowik 2004), which are functions of the desired levels of technical uncertainty which remain unresolved at various points during the innovation process (Van de Ven and Polley 1992, Lenfle and Loch 2010, Azuolay et al. 2011). Therefore, we contribute to the literature which examines the relationship between risk-taking and innovation performance (Staw et al. 1981, March and Shapira 1992, Wiseman and Bromiley 1996), by introducing a contingency model of temporal ambidexterity which considers the effect of a manager's dynamic risk preferences (Das and Teng 1997, Fried and Slowik 2004) on the optimal rates of investment in exploration and exploitation.

### **3.3 A Dynamic Model of Knowledge Creation from Exploration and Exploitation**

We consider a manager leading a team in developing an innovation. During the innovation process, the innovation team's knowledge accumulates in proportion to the investments in knowledge creation that occur over a fixed time horizon  $t \in [0, T]$ , where 0 is the initial time of the innovation process. The terminal time  $T$ , when the innovation process concludes, is given. The technical performance of the innovation is continuously monitored and reflects the changes which result from various knowledge creation activities including concept generation, trial-and-error activities, simulations, prototype

testing and experiments (Van de Ven and Polley 1992, Browning et al. 2002, Austin et al. 2012). The uncertain nature of the technical performance of the innovation under development is captured as the random variable  $X(t)$ , distributed with mean  $\mu(t)$  and variance (standard deviation)  $\sigma(t)$  (March 1991, Huchzermeier and Loch 2001, Carrillo and Gaimon 2004, He and Wong 2004). The *mean*,  $\mu(t)$ , represents the average technical performance of the innovation at time  $t$ . The *variance (standard deviation)*,  $\sigma(t)$ , reflects the projected range and predictability of the technical performance outcomes which could be realized from the innovation at time  $t$  (Huchzermeier and Loch 2001, Browning et al. 2002, Luo et al. 2005).

Let  $\mu(0)>0$  represent the mean technical performance of the innovation at the beginning of the innovation process. Also, let  $\sigma(0)>0$  represent the level of variance of the innovation's technical performance at the beginning of the innovation process. Beyond the innovation team's initial level of knowledge, investments in knowledge creation activities dynamically alter the distribution of the innovation's possible technical performance outcomes (March 1991). The manager directs the innovation team to invest in both exploration and exploitation knowledge creation activities in order to improve the technical performance of the innovation under development. Let the decision variables  $i(t)$  and  $e(t)$  denote the *rate of knowledge creation efforts from exploration* and the *rate of knowledge creation efforts from exploitation* at time  $t \in [0, T]$ , respectively, where  $i(t)$  and  $e(t) \geq 0$ . The mathematical relationships captured in Equations (1) and (2) reflect the dynamic effects of exploration and exploitation on the mean and variance of the distribution of technical performance outcomes. Let  $dY/dy$  denote the first order derivative of  $Y$  with respect to  $y$ . Therefore,  $d\mu(t)/dt$  and  $d\sigma(t)/dt$  denote the rates of

change with respect to time in the mean and variance at time  $t$ . The coefficients  $\alpha_0$  and  $\beta_0 > 0$  in Equation (1) represent the extent to which a unit of exploration versus exploitation positively impacts the mean value of the technical performance outcomes, respectively. The positive coefficients  $\alpha_1$  and  $\beta_1 > 0$  in Equation (2) represent the extent to which a unit of exploration versus exploitation impact the variance of the technical performance outcomes. As demonstrated in Chapter 2, investing in exploration has a variance increasing effect on the range of possible outcomes from the innovation process (Fleming 2001, Austin et al. 2012). In contrast, exploitation reduces the level of technical uncertainty of the innovation outcomes. We model the variance-reducing impact of exploitation in the second term in Equation (2) and the variance-increasing effect of exploration in the first term.

Our model also recognizes the impact of the lagged realization of performance benefits from the knowledge creation process (Gaimon 1997, Carrillo and Gaimon 2000, Rahmandad 2008). March (1991) highlights that characterizations of exploration and exploitation must account for differences not only in the variability, but also in the timing of the performance outcomes from these two modes of knowledge creation; the former being not only “less certain” but also “more remote in time” (p. 73) To account for this feature, our model contrasts the instantaneous impact of exploitation versus the lagged effect of investing in exploration. That is, we assume that the impact from investments in exploration at time  $\tau$ , are not realized until some later time  $t \in [\tau, T]$ , when changes to both the mean and variance due to exploration take effect (Loch and Tapper 2002). We define the distributed lag function  $\theta(\tau-t)$  to reflect the portion of the changes in the mean and variance realized at time  $\tau$ , due to investments in exploration at time  $t$ , such that

$0 \leq \theta(\tau-t) \leq 1$  and  $\int_t^\infty \theta(\tau-t) d\tau = 1$  (Gaimon 1997, Carrillo and Gaimon 2000). The lag function supports a general characterization of when the results from investing in exploration are realized over time. The distributed lag function also allows for the possibility that the results from exploration may not be feasibly realized with the planning horizon given that  $T < \infty$  holds (Loch and Tapper 2002). In the numerical analysis presented in Section 3.6 we characterize the lag using a Gamma distribution (Gaimon 1997, Carrillo and Gaimon 2000).

We also consider the innovation team's *absorptive capacity*, which reflects the team's prior knowledge and cumulative experience, which impacts the efficiency of future exploration and exploitation (Cohen and Levinthal 1990, Rothaermel and Alexandre 2009). First, we consider the effect of absorptive capacity as it relates to the changes in the mean over time. We assume the rates at which both exploration and exploitation improve the mean technical performance at time  $t$ ,  $(d\mu(t)/dt)$ , are increasing in the current level of the mean,  $(\mu(t))$ . That is, we assume that if the innovation team has greater prior knowledge or cumulative experience it is better able to improve the mean level of the innovation outcome at time  $t$ . Mathematically, this relationship is reflected in Equation (1) below.

Next, we consider the effect of the absorptive capacity as it relates to changes in the variance over time  $t$ . Recall that the variable  $\sigma(t)$ , captures the effects of all past investments in exploration and exploitation knowledge creation activities. That is, the innovation outcomes will have a large variance ( $\sigma(t)$  is large) given significant prior investments in exploration. Conversely, a small variance reflects considerable prior investments in exploitation. We assume the rate at which exploration generates variance



$(d\sigma(t)/dt)$  at time  $t$  is increasing in the current level of the variance  $(\sigma(t))$ , as captured in the first term in Equation (2). First, suppose the innovation team has a greater level of cumulative exploration experience (large variance), Equation (2) suggests that it is easier to generate additional novel outcomes from exploration in the future. Empirical research supports this assumption (Lavie and Rosenkopf 2006, Dunlap-Hinkler et al. 2010). In addition, we also assume that the rate at which exploitation reduces the variance is increasing in the current level of the variance  $(\sigma(t))$ , as captured in the second term in Equation (2). Empirical research also supports this assumption that greater levels of cumulative exploration experience (large variance), create more opportunities to refine knowledge and resolve existing technical uncertainty through exploitation in the future (Schilling et al. 2003). Second, suppose the innovation team has a greater level of cumulative exploitation experience, so that the variance is small. As a consequence, the innovation team may suffer from inertia and core rigidities, such that the ability to generate creative outcomes and increase the variance from exploration in the future is more difficult to achieve (Benner and Tushman 2003, Ward 2004). Lastly, we capture the impact of diminishing returns to exploitation. That is, if the innovation team has a greater level of cumulative exploitation experience, so that knowledge is very reliable and well understood (small variance), then there are limits to additional performance refinements, and further reductions in the variance from exploitation are increasingly difficult to achieve (Fleming 2001, Linton and Walsh 2004, Erat and Kavadias 2008). The above discussion is captured in Equations (2). Note that since  $\sigma(0) > 0$ , we have  $\sigma(t) > 0$  holds for  $t \in [0, T]$ .

$$\frac{d\mu(t)}{dt} = \int_0^t \alpha_0 \theta(t-\tau) i(\tau) \mu(\tau) d\tau + \beta_0 e(t) \mu(t) \quad (1)$$

$$\frac{d\sigma(t)}{dt} = \int_0^t \alpha_1 \theta(t-\tau) i(\tau) \sigma(\tau) d\tau - \beta_1 e(t) \sigma(t) \quad (2)$$

### 3.4 Manager's Performance Objectives

#### 3.4.1 Short-term versus Long-term Performance Objectives

Having concluded our discussion of the dynamics of exploration and exploitation we now discuss the manager's innovation performance objectives. The manager's objective in the short-term is to demonstrate ongoing progress in enhancing the technical performance outcomes of the innovation, throughout the planning horizon. Secondly, in the long-term the manager considers the impact of the terminal technical performance on the innovation's future market value, which will be evaluated at the end of the innovation process. In the following discussions, we will denote the manager's short-term and long-term objectives, by  $j=S,L$  respectively.

In the short-term, the innovation process is often monitored, on an ongoing basis, by corporate management, venture capitalists and funding partners (Van de Ven and Polley 1992) who provide tangible or intangible incentives to the manager for demonstrating ongoing performance improvements during the innovation process (Van de Ven and Polley 1992, Chao et al. 2009, Sommer and Loch 2009). As a result of these performance requirements the manager may be driven to be risk-averse or risk-seeking in the short-term. For example, Azoulay et al. (2011) describe the difference between the practices for monitoring progress on NIH (National Institute of Health) versus HHMI (Howard Hughes Medical Institute) innovation grants. The former encourages the researcher to pursue more incremental innovations and requires very detailed and comprehensive preliminary evidence of success, in order to receive ongoing funding. On the other hand, HHMI "urges its researchers to take risks, to explore unproven avenues,

to embrace the unknown even if it means uncertainty or the chance of failure" (Howard Hughes Medical Institute Annual Report 2003, pg. 12).

However, in addition to demonstrating short-term performance improvements, the manager also considers how his decisions may impact achieving his long-term innovation objectives. Naturally, the future market value of the innovation is a function of the potential value from licensing, venture IPO, or commercializing the innovation (Tyebjee, and Bruno 1984, Kerber 2004, Sood and Tellis 2009, Gaimon and Bailey 2012), which increases in relation to its technical performance. For example, Sood and Tellis (2009) empirically demonstrate the impact of terminal outcomes, such as demonstration of a prototype or granting of a patent, on firm market value. However, the market valuation of the innovation also depends on the level of uncertainty associated with realizing the anticipated performance benefits (Luo et al. 2005). Managers exhibit a range of preferences with respect to the variability of payoffs which will be realized from an innovation. A manager who is risk-averse with respect to the market value of the innovation perceives a higher potential payoff when the variance of the innovation's terminal outcomes is small. However, a manager who is risk-seeking in the long-term is willing to develop a more high risk technology innovation, which has a greater upside potential to achieve superior market value, even at the expense of higher downside risk. The literature reflects this characterization of the relationship between risk-taking and long-term payoffs such as stock options (Rajgopal and Shevlin 2002, Xue 2007).

### **3.4.2 Evaluating Innovation Performance**

Our characterization of the manager's performance focus reflects a common distinction in the risk-taking and aspiration literature. This stream of literature

distinguishes between a risk-averse decision-maker who focuses on a survival target, and is therefore mostly concerned with the worst case outcome of the innovation's projected performance, versus a risk-seeking decision-maker who focuses on an aspiration target, and is therefore mostly concerned with the best-case outcome of the innovation's projected performance (Dickson and Giglierano 1986, March and Shapira 1992, Brockner et al. 2004, Boyle and Shapira 2012).

We consider how the manager's decisions to explore and to exploit during the innovation process impact the innovation's projected performance outcomes (Huchzermeier and Loch 2001, Browning et al. 2002, Hillson 2002, Luo et al. 2005). Therefore, our model consider three common performance indicators: (i) the *lower bound (minimum)* of the innovation's projected performance (Baumol 1963, Browning et al. 2002, Hillson 2002, Luo et al. 2005), (ii) the *upper bound (maximum)* of the innovation's projected performance (Browning et al. 2002, Hillson 2002, Luo et al. 2005), or (iii) the *expected value (mean)* of the innovation's projected performance, which is typically assumed to be the focal performance indicator. We leverage and extend the process capability index model (Kane 1986) to characterize the effect of uncertainty on the upper and lower bounds of the innovation's projected performance. Let  $P$  be the standardized normal statistic. Employing the standardized normal transformation  $z_{P_j}(t)$ , as defined in Equation (3), captures the upper (lower) performance confidence limit which is achievable by the innovation at a specified probability, reflected by the standardized normal statistic of  $P_j$  where  $j=S,L$ . The manager's objective is to maximize  $z_{P_j}(t)$ , that is, to maximize the lower (upper) bound of the innovation's projected performance given  $P_j < 0 (> 0)$ . (Baumol (1963) considers a similar formulation for the maximization of the

expected gain-confidence limit, but only for  $P_j < 0$ . Given our intent to focus on the impact of uncertainty, we do not consider the case when  $P_j$  is zero (i.e., the manager's objective is simply to maximize the mean expected performance of the innovation). We consider the implications of the sign of  $P_j$  below.

$$z_{pj}(t) = \mu(t) + P_j \sigma(t), j = S, L \quad (3)$$

With  $P_j < 0$ , the risk-averse manager has a maxi-min objective (Libby and Fishburn 1977). A larger absolute value of  $P_j < 0$  indicates that the manager has a larger degree of risk-aversion, so that he is trying to improve the degree of confidence associated with the innovation's projected minimum performance (Baumol 1963). On the other hand, when  $P_j > 0$ , the manager is risk-seeking. A larger absolute value of  $P_j > 0$  indicates that, while the manager is trying to improve the maximum performance limit of the innovation, he is willing to do so with a higher level of uncertainty and more tolerance for the probability of failure. This is consistent with the need to facilitate risk-taking in order to focus on the innovation's upside potential, and to create breakthrough next-generation innovations (Van de Ven and Polley 1992). Based on this characterization, we refer to  $P_j$  as the *risk-taking preference* with respect to both the short-term and long-term objectives,  $j = S, L$ . Our characterization of the varying risk preferences, as exhibited by venture investors and hence innovation managers, during the innovation process reflect those which have been empirically observed (Lenfle and Loch 2010, Azuolay et al. 2011, Tian and Wang 2011).

Our model also captures the possibility that the manager's short-term and long-term risks preferences are not aligned (Das and Teng 1997). Van de Ven and Polley (1992) observe that managers are often risk-averse in the short-term, with respect to not

being able to demonstrate sufficient progress, during the innovation process, to ensure ongoing funding, but may be risk-seeking with respect to long-term market-related objectives, that is  $P_S > 0$  and  $P_L < 0$  hold. Alternatively, we also consider a manager to be risk-averse in the long-term with respect to market objectives, but risk-seeking in the short-term, with respect to demonstrating development progress, that is  $P_S > 0$  and  $P_L < 0$  hold (Wu and Knott 2006).

### **3.4.3 Cost Objectives**

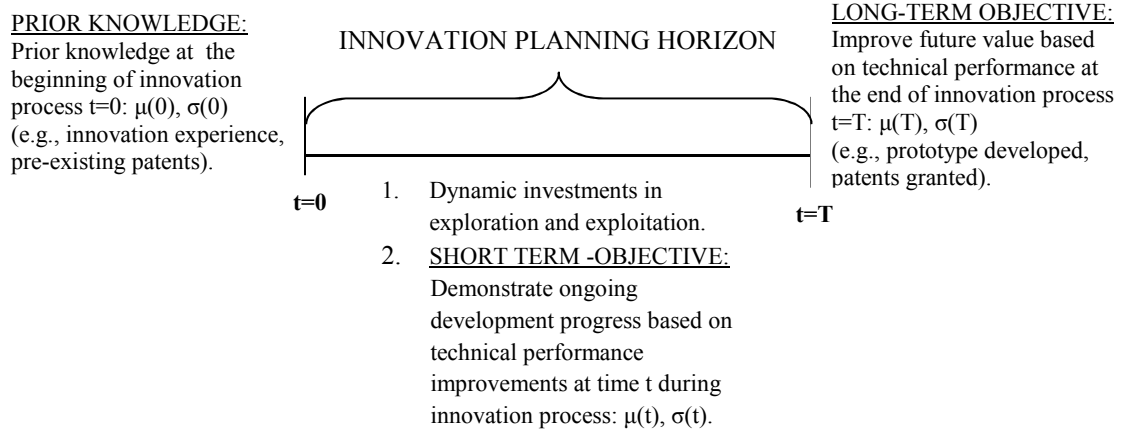
To complete our presentation of the manager's objective function, we consider the impact of the operating costs which are incurred as the firm invests in exploration and exploitation (Bajaj et al. 2004, Choi et al. 2008, Austin et al. 2012, Fixson and Marion 2012). Thomke (1998, 2001) highlights how different cost structures may drive different choices in the optimal mode and sequencing of knowledge creation activities. We define the cost parameters for knowledge creation from exploration and exploitation as  $c_0$  and  $c_1$ , respectively. We make no assumptions about the comparative costs of exploration versus exploitation a-priori. Consistent with the literature, we assume quadratic cost functions for both modes of knowledge creation to reflect the diseconomies of scale due to the disruption and coordination of larger-scale exploration or exploitation activities at any single instant in time (Carrillo and Gaimon 2000).

### **3.4.4 Objective Function**

Figure 2.1 below describes the manager's knowledge creation activities and performance objectives. To summarize, the manager optimally invests in exploration and exploitation to optimize a multi-dimensional objective: (i) maximize the innovation's lower ( $P_S < 0$ ) or upper ( $P_S > 0$ ) performance bound in the short-term during the innovation

process; (ii) maximize the innovation's lower ( $P_L < 0$ ) or upper ( $P_L > 0$ ) terminal performance realized in the long-term at the end of the innovation process; and (iii) minimize the cumulative expenditures incurred for knowledge creation over the innovation process. The marginal incentive for a unit increase in technical performance during the innovation process and at the terminal time of the innovation process are given by the *short-term performance incentive*,  $w_S$ , and the *long-term performance incentive*,  $w_L$ , respectively. The three-part objective is captured in Equation (4), with  $z_{PS}$  and  $z_{PL}$  given as in Equation (3).

$$\text{Maximize } \int_0^T \{w_S z_{PS}(t) - 1/2 c_0 i^2(t) - 1/2 c_1 e^2(t)\} dt + w_L z_{PL}(T) \quad (4)$$



**Figure 3.1: Innovation Process Dynamics and Objectives**

### 3.5 Optimal Dynamic Knowledge Creation Strategies

The manager optimizes Equation (4) above subject to the dynamics in Equations (1) and (2) and non-negativity constraints on  $e(t)$  and  $i(t)$  for all  $t \in [0, T]$ . The model is solved using optimal control theory methods for problems with continuously distributed time lags (Hartl and Sethi 1984, Sethi and Thompson 2000). Throughout the remainder

of the paper "\*" refers to an optimal solution. The optimality conditions and proofs appear in Appendix B.

We introduce the adjoint variables,  $\lambda_0(t)$  and  $\lambda_1(t)$ , to represent the instantaneous marginal values to the objective of a unit increase in the mean and variance at time  $t$ , respectively. We also introduce the variables,  $x_0(t)$  and  $x_1(t)$ , to represent the cumulative distributed marginal values of a unit increase in the mean and variance at time  $t$ , respectively, due to the lagged effect of exploration, given in Equations (5) and (6) below. Theorem 1 and Corollaries 1a and 1b illustrate how the innovation manager values a unit increase in the mean and the variance at time  $t$ , throughout the planning horizon.

$$x_0(t) = \int_t^T \theta(\tau - t) \lambda_0(\tau) d\tau \quad (5)$$

$$x_1(t) = \int_t^T \theta(\tau - t) \lambda_1(\tau) d\tau \quad (6)$$

**THEOREM 1.** *The marginal value of a unit increase in the mean and variance at time  $t$  satisfy (i) and (ii), respectively:*

$$(i) \frac{d\lambda_0}{dt} = -w_s - \alpha_0 i(t) x_0(t) - \lambda_0(t) \beta_0 e(t) \text{ and } \lambda_0(T) = w_L > 0$$

$$(ii) \frac{d\lambda_1}{dt} = -w_s P_s - \alpha_1 i(t) x_1(t) + \lambda_1(t) \beta_1 e(t) \text{ and } \lambda_1(T) = w_L P_L$$

**COROLLARY 1a.** *The instantaneous and cumulative distributed marginal values of a unit increase in the mean ( $\lambda_0(t)$ ,  $x_0(t)$ ) are both positive and non-increasing for  $t \in [0, T]$ .*

**COROLLARY 1b.** *The instantaneous and cumulative distributed marginal value of a unit increase in the variance ( $\lambda_1(t)$ ,  $x_1(t)$ ) may be positive or negative, and non-decreasing or non-increasing. Case (I)-(IV) (below) depict four scenarios that are likely to occur for*



$t \in [0, T]$ . We introduce conditions corresponding to all possible solutions for  $\lambda_I(t)$  and  $x_I(t)$  in Appendix B.

(I) Given  $P_S < 0$  and  $P_L < 0$  hold, then  $\lambda_I(t) < 0$ ;  $d\lambda_I(t)/dt > 0$ ;  $x_I(t) > \lambda_I(t)$ ,  $dx_I(t)/dt > 0$  are likely to occur for  $t \in [0, T]$ ,

(II) Given  $P_S > 0$  and  $P_L > 0$  hold, then  $\lambda_I(t) > 0$ ;  $d\lambda_I(t)/dt < 0$ ;  $x_I(t) < \lambda_I(t)$ ,  $dx_I(t)/dt < 0$  are likely to occur for  $t \in [0, T]$ ,

(III) Given  $P_S > 0$  and  $P_L < 0$  hold, then  $\lambda_I(t) > 0$ ,  $d\lambda_I(t)/dt < 0$ ,  $t \in [0, t_1]$  and  $\lambda_I(t) < 0$ ,  $d\lambda_I(t)/dt < 0$ ,  $t \in [t_1, T]$  are likely to occur for  $t \in [0, T]$ ,

(IV) Given  $P_S < 0$  and  $P_L > 0$  hold, then  $\lambda_I(t) < 0$ ,  $d\lambda_I(t)/dt > 0$ ,  $t \in [0, t_1]$  and  $\lambda_I(t) > 0$ ,  $d\lambda_I(t)/dt > 0$ ,  $t \in [t_1, T]$  are likely to occur for  $t \in [0, T]$ .

From Corollary 1a, the marginal value of an unit increase in the mean performance is always positive and non-increasing, so that it is always more valuable to invest in an additional unit of the mean technical performance earlier in the development horizon in order to maximize the remaining time to leverage the mean absorptive capacity benefits. Corollary 1b indicates that the marginal value of an additional unit of variance may be either positive or negative, and that increasing or decreasing the variance may be more valuable earlier or later in the development horizon, depending on the innovation team's short-term versus long-term performance objectives.

From Corollary 1b(I), when the manager is risk-averse in both the short-term and the long-term ( $P_S < 0$  and  $P_L < 0$ ), the instantaneous marginal value of the variance is negative and increasing. Under these conditions, the magnitude of the cumulative distributed marginal value of the variance is large relative to the instantaneous marginal value of the variance ( $x_I(t) \geq \lambda_I(t)$ ). That is, when resolving uncertainty is valuable

( $\lambda_1(t) < 0$ ), the delay in realizing the performance impacts of exploration increases the risk-averse manager's propensity to invest in variance-generating exploration activities. This is an interesting and counterintuitive finding. However, these insights are consistent with empirical findings on the impact of time delays on risk-taking decisions. For example, in an experimental study Abdellaoui et al. (2011) find that participants become increasingly more risk-tolerant with delayed lotteries. On the other hand, from Corollary 1b(II), when the manager is risk-seeking in both the short-term and the long-term ( $P_S > 0$  and  $P_L > 0$ ), the instantaneous marginal value of the variance is positive and decreasing. Under these conditions, the magnitude of the cumulative distributed marginal value of the variance is small relative to the instantaneous marginal value of the variance ( $x_1(t) \leq \lambda_1(t)$ ). That is, when generating uncertainty is valuable ( $\lambda_1(t) > 0$ ), the delay in realizing the performance impacts of exploration reduces the risk-seeking manager's propensity to invest in exploration.

Further interpretation of the marginal values of the mean and variance are given when we discuss the optimal dynamic exploration and exploitation strategies in Theorems 2 and 3 which follow. Theorem 2 provides insights on how the marginal value of the mean and variance, along with the innovation team's level of knowledge and knowledge creation capabilities affects the manager's decision to either increase or decrease the relative focus on investing in exploration versus exploitation.

**THEOREM 2.** *The manager optimally invests in exploitation and exploration for  $t \in [0, T]$ :*

$$(i) \ e^*(t) = \text{Max} \left\{ \frac{\lambda_0(t)\beta_0\mu(t) - \lambda_1(t)\beta_1\sigma(t)}{c_1}, 0 \right\}; \quad (ii) \ i^*(t) = \text{Max} \left\{ \frac{x_0(t)\alpha_0\mu(t) + x_1(t)\alpha_1\sigma(t)}{c_0}, 0 \right\}.$$

**COROLLARY 2a.** *There may exist an optimal stopping (starting) time for exploration or exploitation, such that after (before) the optimal stopping (starting) time, it is not optimal to invest in knowledge creation.*

**COROLLARY 2b.** *Given  $e^*(t) > 0$ , the rate of exploitation is larger if: (i)  $\mu(t)$  or  $\lambda_0(t) > 0$  is larger; (ii)  $\sigma(t)$  is larger (smaller) when  $\lambda_1(t) < 0$  ( $\lambda_1(t) > 0$ ); or (iii)  $c_1$  is smaller.*

**COROLLARY 2c.** *Given  $i^*(t) > 0$ , the rate of exploration is larger if: (i)  $\mu(t)$  or  $x_0(t) > 0$  is larger; (ii)  $\sigma(t)$  is smaller (larger) when  $x_1(t) < 0$  ( $x_1(t) > 0$ ); or (iii)  $c_0$  is smaller.*

The analytical results of Theorem 2 and Corollary 2a provide insights on the tensions faced by managers when implementing an ambidextrous knowledge creation strategy. First, both exploration and exploitation improve mean performance, as noted in Corollaries 2b and 2c. Second, as the marginal value of an additional unit of variance ( $\lambda_1(t)$ ,  $x_1(t)$ ) becomes increasingly positive (negative), then the optimal rate of exploitation decreases (increases), while the optimal rate of exploration increases (decreases). Third, from Equation (2) while the variance is decreasing in the rate of exploitation, the marginal effectiveness of exploitation for reducing the variance is improved at higher levels of the variance, which results from earlier exploration. Our model therefore captures the paradoxical challenges associated with balancing exploration and exploitation. Exploration and exploitation are simultaneously (i) substitutes in improving the mean, (ii) opposites with respect to increasing versus decreasing the variance and (iii) complements in the long-term benefits of higher levels of the variance, as a result of prior exploration, which enhances the effectiveness of future exploitation.

March (1991) cautions of the potential dangers of over-exploring and over-exploiting. However, as yet, there are no clear guidelines for determining at what point exploration and exploitation transition from being helpful to being harmful. From Theorem 2(i) and Corollary 2a, if the instantaneous marginal value of the variance is negative ( $\lambda_1(t) < 0$  for all  $t \in [0, T]$ ), the manager optimally undertakes exploitation for the entire innovation process. However, if the marginal value of the variance is positive ( $\lambda_1(t) > 0$  for  $t \in [0, T]$ ), there may exist an optimal stopping (starting) point for exploitation. That is, if  $\lambda_1(t)\beta_1\sigma(t) > \beta_0\lambda_0(t)\mu(t)$ , so that the instantaneous marginal value of the variance at time  $t$  is positive, but either the exploitation variance-reducing capability or the current variance is also large, then at some time  $t \in [0, T]$  during the planning horizon it may no longer be optimal for the manager to invest in exploitation. On the other hand, From Theorem 2(i) and Corollary 2a, if the cumulative distributed marginal value of the variance is positive ( $x_1(t) > 0$  for all  $t \in [0, T]$ ), then the manager optimally undertakes exploration for the entire duration of the innovation process. However, if  $x_1(t) < 0$  where  $-x_1(t)\alpha_1\sigma(t) > \alpha_0x_0(t)\mu(t)$ , then at some time  $t \in [0, T]$  it may not be optimal to explore if the cumulative distributed marginal value of the variance is negative, but exploration increases the variance at a higher rate than it increases the mean. While the notion of exploration and exploitation learning traps are conceptually well-accepted (Levitt and March 1988, Levinthal and March 1993), our model provides an important new analytical perspective for identifying the optimal stopping (starting) point in order to avoid over-investing in exploration and exploitation.

The above discussions assume that both exploration and exploitation on average increase the mean performance ( $\alpha_0, \alpha_1 > 0$ ). The importance of this assumption is clear.

Given a scenario in which the marginal value of the variance is negative, then if exploration also decreases the mean expected performance ( $\alpha_0 < 0$ ), Equation (ii) in Theorem 2 is negative and it is never optimal to explore. As such, a feasible balancing strategy between exploration and exploitation would not exist, thereby leading to the trivial solution of investing in exploitation only. Although we do not explicitly consider the case where exploration decreases the mean performance ( $\alpha_0 < 0$ ), we do however consider scenarios where exploration has negligible benefits for improving the innovation's mean performance ( $\alpha_0 \approx 0$ ).

### 3.5.1 Sequential versus Fixed-Dominant Knowledge Creation Strategies

Given the dynamic nature of the knowledge creation process, changes in the marginal value of the variance can cause the optimal relative emphasis on investments in exploration and exploitation to change over time (Thomke 1998). As a result, following from Theorems 2 and Corollary 3 below, we are able define four possible dynamic strategies which may occur, which are highlighted in Table 3.1.

**Table 3.1: Possible Temporal Ambidexterity Strategies**

Fixed Dominance	Fixed Exploitation–Dominant	Exploitation exceeds exploration throughout innovation process.
	Fixed Exploration–Dominant	Exploration exceeds exploitation throughout innovation process.
Sequential Dominance	Sequential Explore–Exploit	Exploration initially exceeds exploitation, then vice versa.
	Sequential Exploit–Explore	Exploitation initially exceeds exploration, then vice versa.

**COROLLARY 3.** *An optimal switching time may exist at  $t_s$  where  $t_s \in (0, T)$  such that  $i^*(t_s) = e^*(t_s)$  holds. In particular, if  $i^*(t_s) > e^*(t_s)$  for  $t < t_s$  and  $i^*(t_s) < e^*(t_s)$  for  $t > t_s$ , the manager optimally pursues a sequential explore-exploit strategy. If the reverse holds, the manager optimally pursues a sequential exploit-explore strategy.*

The manager may optimally pursue a *fixed exploration-dominant* or *fixed exploitation-dominant* strategy if the conditions  $i^*(t) > e^*(t)$  or  $e^*(t) > i^*(t)$  holds for all  $t \in [0, T]$ , respectively. In the fixed exploitation-dominant strategy the rate of exploitation always exceeds the rate of exploration, so that uncertainty is continuously reduced over time. This strategy is consistent with a stage-gate innovation process, which is focused on screening an initial population of ideas for feasibility and continuously reducing uncertainty over time (Cooper 1990). In contrast, in the fixed exploration-dominant strategy the rate of exploration always exceeds the rate of exploitation, so that uncertainty is increasing over time. This suggests a process of ongoing search for new ideas.

As the process of knowledge creation proceeds, the marginal benefit of one or the other mode of knowledge creation could either increase or decrease over time. Therefore, the manager may optimally switch the team's focus from one dominant mode of knowledge creation to the other at a switching time,  $t_s$  as noted in Corollary 3 (Thomke 1998). We refer to the dynamic strategy as a *sequential explore-exploit (sequential exploit-explore)* strategy if before the switching time the rate of exploration (exploitation) optimally dominates, and after the switching time the rate of exploitation (exploration) optimally dominates.

A *sequential explore-exploit* sequence suggests that the manager initially focuses on increasing the variance and then switches to focus on resolving uncertainty later in the innovation process. This sequential strategy is the most commonly cited strategy in the literature (Rothaermel and Deeds 2004, Terwiesch and Ulrich 2009). However, from Corollary 3, we are also able to identify conditions under which a sequential exploit-explore dynamic knowledge creation strategy optimally occurs, in which the innovation

team initially focuses on resolving uncertainty and later focuses on generating new ideas (Koput 1997, Thomke 1998, Biazzo 2009). Therefore, our model provides analytic conditions under which this atypical exploit-explore dynamic knowledge creation strategy could arise. Further discussion of the optimal timing of the peak rates of exploration and exploitation is given in the next section.

### 3.5.2 Front-Loaded versus Back-Loaded Knowledge Creation Strategies

Beyond determining the relative rates of exploration and exploitation, we are also interested in defining at what point during the innovation process the innovation team optimally focuses the peak efforts on each activity. We define the knowledge creation efforts as *front-loaded* when the rate of a knowledge creation activity is positive and decreasing ( $de^*(t)/dt < 0$  or  $di^*(t)/dt < 0$  for all  $t \in [0, T]$ ), whereby the peak rate of investment occurs at the initial time (Thomke and Fujimoto 2000, Thomke 2001, Carrillo and Gaimon 2004, Ozkan et al. 2009). There are various theories which promote the optimality of front-loading development activities. Cohen and Levinthal's (1990) notion of absorptive capacity suggests front-loading knowledge creation investments in order to maximize the benefits of cumulative knowledge, which improves the marginal effectiveness of future knowledge creation activities. In contrast, we define the knowledge creation efforts as *back-loaded* when the rate of a knowledge creation activity is positive and increasing ( $de^*(t)/dt > 0$  or  $di^*(t)/dt > 0$  for all  $t \in [0, T]$ ), whereby the peak rate of investment occurs at the terminal time (Carrillo and Gaimon 2004, Ozkan et al. 2009, Fixson and Marion 2012). Finally, we define the knowledge creation strategy as a *delay* strategy if the peak rate of investments optimally occurs at some intermediate time  $t \in (0, T)$ , whereby we obtain an inverse U-shaped solution.

**THEOREM 3.** *The optimal rates of change of investment in exploitation and exploration at time  $t$  satisfy the following for  $e^*(t) > 0$  and  $i^*(t) > 0$  and  $t \in [0, T]$ :*

$$(i) \quad \frac{de(t)^*}{dt} = \frac{\beta_0 \left[ \lambda_0(t) \frac{d\mu(t)}{dt} + \frac{d\lambda_0(t)}{dt} \mu(t) \right] - \beta_1 \left[ \lambda_1(t) \frac{d\sigma(t)}{dt} + \frac{d\lambda_1(t)}{dt} \sigma(t) \right]}{c_1}$$

$$(ii) \quad \frac{di(t)^*}{dt} = \frac{\alpha_0 \left[ x_0(t) \frac{d\mu(t)}{dt} + \frac{dx_0(t)}{dt} \mu(t) \right] + \alpha_1 \left[ x_1(t) \frac{d\sigma(t)}{dt} + \frac{dx_1(t)}{dt} \sigma(t) \right]}{c_0}$$

Taken together, Theorem 3 along with Corollaries 1(a) and 1(b) provide the conditions under which the manager optimally front-loads, delays or back-loads the peak levels of investment in exploitation and exploitation. As a result of the benefits of the mean to the objective, both exploration and exploitation are driven to be front-loaded. Additionally, for both exploration and exploitation, the two terms which are multiplied by  $\alpha_1$  and  $\beta_1$ , could be either positive or negative. Given  $\sigma(t)$  is always positive, then the determination of whether the manager invests in exploration and exploitation at an increasing or decreasing rate is a function of: (i) the instantaneous and cumulative distributed marginal value of the of variance ( $\lambda_1(t)$ ,  $x_1(t)$ ); (ii) the marginal value of generating versus reducing the variance, earlier versus later in the innovation process ( $d\lambda_1(t)/dt$ ,  $dx_1(t)/dt$ ); and (iii) the innovation team's relative rates of variance reduction versus variance at time  $t$ , which determine the sign of  $d\sigma(t)/dt$ . Based on these three drivers, we discuss conditions under which the peak investments in exploitation and exploration are driven to occur earlier versus later in the innovation process.



**COROLLARY 4a.** *The manager optimally invests in exploitation at a non-increasing rate at time  $t$ , if the following conditions hold for  $t \in [0, T]$ : (i)  $\lambda_1(t) < 0$ ,  $d\lambda_1(t)/dt > 0$ ,  $d\sigma(t)/dt < 0$  or (ii)  $\lambda_1(t) > 0$ ,  $d\lambda_1(t)/dt > 0$ ,  $d\sigma(t)/dt > 0$ .*

If the conditions in Corollary 4a(i) hold then exploitation is front-loaded ( $de^*(t)/dt < 0$ ). Based on Corollary 1b.I when the manager's objectives are aligned, such that he risk-averse in both the short-term and the long-term ( $P_s < 0$  and  $P_L < 0$ ), then it is likely that the marginal value of the variance is negative and increasing ( $\lambda_1(t) < 0$ ,  $d\lambda_1(t)/dt > 0$ ). Also, consider a situation where the variance would be decreasing over time ( $d\sigma(t)/dt < 0$ ). This situation occurs when the team's variance reduction capabilities outpace the rate of variance generation. For example, the innovation team may demonstrate superior variance reduction capabilities if they have access to advanced technologies for testing, screening and other uncertainty reduction methods, that is if  $\beta_1 \gg \alpha_1$  (Thomke 2001). Under the conditions defined in Corollaries 4a(i) and 1b.I, investing in exploitation to reduce the variance is not only more valuable, but also more effective, earlier in the innovation process, as the marginal variance reduction from exploitation is highest when variance is at its maximum at the beginning of the innovation process. Alternatively, from Corollaries 4a(ii) exploitation is front-loaded if the marginal value of the of variance is positive and increasing for all  $t \in [0, T]$ , ( $\lambda_1(t) > 0$ ,  $d\lambda_1(t)/dt > 0$ ) while the variance is increasing ( $d\sigma(t)/dt > 0$ ). However, from Theorem 2, since the optimal rate of exploitation decreases as the marginal value of the variance becomes positive this second scenario is unlikely to occur, unless the mean effects from exploitation dominate the variance effects.

Recall from Theorem 3(i) that the sum of the two terms which are multiplied by  $\beta_0$ , respectively, are negative, while the two terms which are multiplied  $\beta_1$ , could be either positive or negative. Therefore, from Theorem 3 and Corollary 4a(ii), exploitation is driven towards a back-loaded strategy if the marginal value of the variance is decreasing throughout the innovation process for  $t \in [0, T]$ . From Corollary 1b.III, this scenario is most likely to occur if the short-term and long-term performance objectives are not aligned so that the manager is risk-seeking in the short-term,  $P_S > 0$ , but risk-averse in the long-term,  $P_L < 0$ . A back-loaded strategy is reinforced with increasing variance, so that the variance reduction benefits from absorptive capacity, peak later in the innovation process when the level of the variance reaches its maximum value.

**COROLLARY 4b.** *The manager optimally invests in exploration at a non-increasing rate at time  $t$ , if the following conditions hold for  $t \in [0, T]$ : (i)  $x_1(t) > 0$ ,  $dx_1(t)/dt < 0$ ,  $d\sigma(t)/dt > 0$  or (ii)  $x_1(t) < 0$ ,  $dx_1(t)/dt < 0$ ,  $d\sigma(t)/dt < 0$ .*

If the conditions in Corollary 4b(i) hold then exploration is front-loaded ( $di^*(t)/dt < 0$ ). Based on Corollary 1b.II when the manager's objectives are aligned, such that he risk-seeking in both the short-term and the long-term ( $P_S < 0$  and  $P_L < 0$ ), then it is likely that the cumulative distributed marginal value of the variance is positive and decreasing ( $x_1(t) < 0$ ,  $dx_1(t)/dt > 0$ ), while the variance is increasing ( $d\sigma(t)/dt > 0$ ). From Corollary 4b(ii), exploration is also optimally front-loaded when, among other conditions, the cumulative distributed marginal value of the variance is negative and decreasing over time. However, from Theorem 2, since the optimal rate of exploration decreases as the marginal value of the variance becomes negative this second scenario is unlikely to occur, unless the mean effects from exploration dominate the variance effects.

Recall from Theorem 3(i) that sum of the two terms which are multiplied by  $\alpha_0$ , are negative, while the two terms which are multiplied  $\alpha_1$ , could be either positive or negative. Given  $\sigma(t)$  is always positive, then based on Theorem 3 and Corollary 4b, exploration is more desirable later in the innovation process when the marginal value of the variance is increasing. From Corollary 1b.IV, this scenario is most likely to occur if the short-term and long-term performance objectives are not aligned, so that the manager wants to reduce variance in performance in the short-term,  $P_S < 0$ , and increase variance in the long-term,  $P_L > 0$ . Furthermore, from Theorem 3(ii) a back-loaded exploration strategy is reinforced with increasing variance, so that absorptive capacity peaks later in the development process. The variance is increasing over time when the variance generation capabilities of exploration dominate the variance reducing effects of exploitation ( $\alpha_1 \gg \beta_1$ ). For example, in the drug development process, advances in combinatorial chemistry have facilitated superior variance generation capabilities which allow the automated generation of numerous variants of potential drug compounds (Thomke 2001).

### **3.6 Numerical Analysis and Managerial Insights**

In the analytical insights presented above, we examine factors which determine how various optimal temporal ambidexterity strategies, for simultaneously investing in exploration and exploitation over time may endogenously arise. From Theorem 3, three possible optimal dynamic strategies are identified for each mode of learning: a front-loaded strategy, a back-loaded strategy and a delay strategy. As such, there are nine possible combined temporal ambidexterity strategies. In this section, we present and discuss the results from extensive numerical analysis to illustrate several of the nine possible temporal ambidexterity strategies which may arise, including those not

previously discussed in the literature (see Table 3.2 below).

**Table 3.2: Optimal Temporal Ambidexterity Scenarios**

CASE	SHORT TERM OBJECTIVE	LONG TERM OBJECTIVE	KEY SCENARIO CONDITIONS	OPTIMAL TEMPORAL AMBIDEXTERITY STRATEGY
1A	RISK AVERSE  Maximize Lower Performance Bound	RISK AVERSE  Maximize Lower Performance Bound	Exploration provides limited improvements in mean performance (i.e., $\alpha_0$ small).	<b>Fixed Exploitation-Dominant</b> <i>Front-load Exploitation</i> <i>Exploration does not occur</i>
1B			Exploration provides superior improvements in mean performance (i.e., $\alpha_0 > \beta_0$ ).	<b>Fixed Exploitation-Dominant</b> <i>Front-load Exploitation</i> <i>Delay peak rate of Exploration (Inverse U) with Optimal Starting Point</i>
1C			Small initial variance (i.e., $\sigma(0)$ small)	<b>Fixed Exploitation-Dominant</b> <i>Front-load Exploitation</i> <i>Front-load Exploration</i>
1D			Small initial variance (i.e., $\sigma(0)$ small) Low cost of exploration (i.e., $c_0$ small)	<b>Sequential Explore-Exploit</b> <i>Front-load Exploitation</i> <i>Front-load Exploration</i>
1E			Impact of lag more pronounced so that performance changes resulting from exploration increasingly delayed	<i>Magnitude of exploration increases and peak rate of exploration expedited as lag effect magnified</i>
2A	RISK SEEKING	RISK SEEKING	Impact of lag more pronounced so that performance changes resulting from exploration increasingly delayed	<i>Magnitude of exploration decreases and peak rate of exploration expedited as lag effect magnified</i>
2B	Maximize Upper Performance Bound	Maximize Upper Performance Bound	Small initial variance (i.e., $\sigma(0)$ small)	<b>Fixed Exploitation-Dominant</b> <i>Front-load Exploitation</i> <i>Front-load Exploration</i>
3	RISK SEEKING	RISK AVERSE	Objectives not aligned. Short-term Risk Seeking and Long-term Risk Averse.	<b>Sequential Explore-Exploit</b> <i>Front-load Exploration with Optimal Stopping Point</i> <i>Back-load Exploitation</i>
4	RISK AVERSE	RISK SEEKING	Objectives not aligned. Short-term Risk Averse and Long-term Risk Seeking.	<b>Sequential Exploit-Explore</b> <i>Front-load Exploitation</i> <i>Delay peak rate of Exploration (Inverse U) with Optimal Starting Point</i>

In Cases 1-4 we examine the impact of several key factors which impact the optimal temporal ambidexterity strategy, namely: (i) the initial knowledge with which the innovation team is endowed,  $(\mu(0), \sigma(0))$ , (ii) the marginal effectiveness of exploration and exploitation,  $(\alpha_0, \beta_0, \alpha_1, \beta_1)$ , and (iii) the manager's dynamic short-term and long-

term risk preferences ( $P_S$ ,  $P_L$ ). Details of the numerical analyses are provided in Appendix B.

In Cases 1A, 1B, 1C and 1D, we consider a manager whose short-term and long-term performance objectives are aligned, such that he is risk-averse with respect to performance outcomes achieved in both timeframes (Corollary 1b.I). On the other hand, in Cases 2A and 2B we assume that the manager is risk-seeking in both the short-term and the long-term (Corollary 1b.II). In contrast to Cases 1 and 2, in which the manager's short-term and long-term objectives are aligned, in Case 3 we assume that the manager is risk-seeking in the short-term, (i.e., with an aim to maximize the upper bound during the innovation process), but risk-averse in the long-term (i.e., tries to limit low performance of the final innovation outcomes) (Corollary 1b.III). Alternatively, in Case 4, we assume that the manager is risk-averse in the short-term but risk-seeking in the long-term (Corollary 1b.IV).

In the Cases 1A-1D, the manager is risk-averse in the short-term and long-term. These cases are important since the literature reflects many instances of venture investors and managers who are risk-averse with respect to the performance outcomes both during as well as at the end of the innovation process (Tian and Wang 2011). To explore the temporal ambidexterity strategies which may optimally occur for the risk-averse manager, we consider five subcases.

In Case 1A, we assume that the innovation team begins the innovation process with a large pool of initial ideas, so that the initial variance of the performance outcomes is large (i.e.,  $\sigma(0)$  large). Furthermore, the manager is also able to explore various new innovation alternatives, with the possibility of discovering either a highly effective or

highly ineffective solution (i.e.,  $\alpha_1$  large). However, on average, the exploration provides negligible improvements in the mean performance (i.e.,  $\alpha_0$  small). Under these conditions from Theorem 2(ii), it is not optimal to explore and the manager pursues a front-loaded exploitation-only strategy. The *fixed exploitation-dominant* strategy, which endogenously arises in the risk-averse setting, is related to the stage-gate process, whereby performance uncertainty is continuously reduced over time (see Appendix Figure B5).

In Case 1B, we also assume that the initial variance of the performance outcomes is large (i.e.,  $\sigma(0)$  large). However, we assume that exploration techniques are available with superior capabilities to improve the mean performance (i.e.,  $\alpha_0 \gg \beta_0$ ). Under these conditions the manager optimally front-loads exploitation and delays the peak rate of exploration. However, in this *fixed exploitation-dominant* strategy, exploration is always pursued at a smaller rate and there may also exist an optimal starting time (see Appendix Figure B6). Pursuing exploration in a risk-averse environment is counterintuitive, and this highlights the importance of considering, not only the variance effects, but also the mean effects of exploration on performance.

In Case 1C, consistent with Case 1B, we assume that the innovation team has access to exploration techniques with the capability to improve the mean performance. However, in Case 1C we also assume that the team starts the development process with a low variance, such that the initial range of performance outcomes is very limited and precisely achievable (i.e.,  $\sigma(0)$  small). Under these conditions, it is still optimal for the manager to pursue a *fixed exploitation-dominant* strategy. However, relative to Case 1B (see Appendix, Figure B6), when the initial variance,  $\sigma(0)$ , is small, it is increasingly valuable to invest in exploration (see Appendix, Figure B7). This case highlights the

importance of early exploration, even in a risk-averse environment. When the initial variance is small, the innovation team should optimally engage in early exploration. By adhering to this optimal policy, the team can avoid the inherent pitfall of converging on a final solution, too early (Leonardi 2011). This allows the innovation team to gain a better understanding of the design space and the potential technical alternatives. By generating variance early in the planning horizon, the team is able to observe a wide range of outcomes, including failures, which enhance the marginal effectiveness of future exploitation, in order to meet the risk-averse objectives (Thomke 1998, Thomke and Fujimoto 2000).

In Case 1D, we demonstrate that the propensity to invest in exploration in a risk-averse setting, given low initial variance, is reinforced and magnified as the cost of exploration is reduced (Austin et al. 2012). Compared to Case 1C (see Appendix Figure B7), as the cost of exploration is reduced, the risk-averse manager eventually pursues a *sequential explore-exploit* strategy in which exploration dominates early in the innovation process, and exploitation dominates later (see Appendix Figure B8). Importantly, therefore, we provide a set of conditions, with respect to short and long-term performance objectives, as well as prior knowledge and the cost of learning, under which the managerial policy to “fail-early, fail-cheap” could endogenously arise (Thomke 2001, Austin et al. 2012). Highlighting the benefits of a *sequential explore-exploit* strategy within an innovation project, Lenfle and Loch (2010) qualitatively capture these findings: “*the project manager found himself iteratively working through the process steps... adding modifications that had not been foreseen in the original design...In the end, the project manager was demoted and a new team brought in. This (new) team ... identified*

130 “quality problems,” and fixed these over another 12 months with a strict phased planning approach. After this period, ramp-up was successful, and the facility reached its design capacity. The company concluded that the second team had won the day with better, more disciplined methods. However, this analysis failed to recognize that...the subsequent “rigorous planning” phase succeeded only because this fundamental work had been carried out first.” (p. 13).

In Case 1E, we consider the impact of the lagged realization of the benefits from investing in exploration. As described in Corollary 1b.I, we show that with short-term and long-term risk-averse objectives ( $P_s < 0$  and  $P_L < 0$ ), as the severity of the lagged realization becomes more delayed the cumulative distributed marginal value of the variance becomes larger relative to the instantaneous marginal value of the variance, so that the manager optimally undertakes even greater amounts of exploration early in the development process (see Appendix Figure B9). This scenario numerically demonstrates the interesting and counterintuitive experimental findings presented by Abdellaoui et al. (2011) which empirically demonstrate that performance delays can increase the propensity for risk-taking.

In contrast to the risk-averse setting in Cases 1A-1E, we now consider a manager who is risk-seeking with respect to both the short-term and long-term objectives, such that the marginal value of the variance is positive and decreasing (Corollary 1b.II). In Case 2A, we consider the impact of the lagged realization of the benefits from investing in exploration. In this risk-seeking setting, as the effects of exploration are increasingly delayed, the manager optimally expedites and also reduces the rates of investment in exploration, so that the optimal solution moves from a *fixed dominant-exploration* to a



*sequential explore-exploit* strategy. As described in Corollary 1b.II, as the severity of the lagged realization is more delayed, the cumulative distributed marginal value of the variance is smaller, so that the manager optimally undertakes less exploration (see Appendix Figure B10). This scenario numerically demonstrates the intuitive result that, as the lag becomes more pronounced the incentive to invest in exploration is reduced, with risk-seeking objectives. Our results highlight an important finding that, for managers faced with inter-temporal risk-taking decisions, a time lag between initiating the investment and realizing the outcome from exploration may act as either an incentive or disincentive to invest, depending on whether the manager is risk-averse versus risk-seeking, respectively (Abdellaoui et. al 2011).

In Case 2B, we consider a situation in which the marginal impact to the mean from exploitation is low and the initial variance is small. Under these conditions, we find that the manager optimally pursues a *fixed exploitation-dominant* strategy (see Appendix Figure B11). This is an interesting result, as it specifies conditions under which, counter-intuitively, a risk-seeking manager focuses on exploitation and uncertainty reduction. Managerially, we show that if prior knowledge indicates that the initial performance outcomes have very low variance, the manager optimally pursues low cost incremental innovation activities. In other research, managers are cautioned against an over-reliance on uncertainty-reducing exploitation in a risk-seeking setting, whereby exploitative core capabilities may become core-rigidities (Leonard-Barton 1992). However, our findings provide conditions, under which excessive prior exploitation experience (i.e., small initial variance) may optimally give rise to a continued focus on exploitation, even with risk-seeking objectives.

Lastly, we consider two important scenarios in which short-term versus long-term performance objectives are not aligned (Das and Teng 1997). In Case 3, we consider a manager who is risk-seeking in the short-term but risk-averse in the long term (Wu and Knott 2006). Consistent with Corollaries 1b.III, 4a and 4b the manager optimally pursues a *sequential explore-exploit* strategy, in which he front-loads exploration but back-loads exploitation. Importantly, reflecting the insights from Corollary 2a, there exists an optimal stopping point for exploration (see Appendix Figure B12).

In Case 4, we consider a manager who is risk-averse with respect to his short-term objective but risk-seeking with respect to his long-term objective (Van de Ven and Polley 1992). This scenario may reflect a situation in which the manager seeks to demonstrate and report his intermediate progress with a high degree of certainty in order to ensure sufficient ongoing funding so as not to “sink the boat” (p. 61) (Dickson and Giglierano 1986). However, at the same time, he may be risk-seeking in the longer term, so as not to “miss the boat” (p. 62) regarding the market opportunity for the innovation (Dickson and Giglierano 1986). Consistent with Corollaries 1b.IV, 4a and 4b, the manager optimally front-loads exploitation to reduce uncertainty in the short-term but back-loads exploration toward the end of the innovation process to increase the upside potential of realizing a breakthrough innovation outcome. Note that, there may exist an optimal starting point for exploration. Furthermore, although the manager is risk-seeking in the long term, the peak rate of exploration occurs prior to the terminal time due to the delayed realization of the benefits of exploration (see Appendix Figure B13).

### **3.7 Conclusions**

Within a single organizational unit, pursuing a temporal ambidexterity strategy

has been advocated as an alternative to pursuing an organizational or functional-domain ambidexterity strategy, in which paradoxical exploration and exploitation activities are disbursed across multiple organizational units (Brown and Eisenhardt 1997, Siggelkow and Levinthal 2003, Lavie et al. 2010). There have been several calls to add “time” as an important research lens through which to examine the concept of ambidexterity (Raisch et al. 2009, Lavie et al. 2010, O’Reilly and Tushman 2013). To date, a limited body of research has leveraged qualitative and simulation-based studies to gain insights on the drivers, mechanisms and outcomes related to a temporal ambidexterity strategy. However, a normative model of temporal ambidexterity has not yet been proposed, so that the optimal sequence of exploration and exploitation activities is not specified in the literature. We respond to this gap in the literature by introducing a dynamic model of knowledge creation from exploration and exploitation. We contribute to the theory of temporal ambidexterity by providing new analytical insights on the inter-temporal tradeoffs, conflicts and synergies which may occur when simultaneously pursuing exploration and exploitation.

Our model embodies several elements that have been shown to be empirically and conceptually critical in theorizing a model of ambidexterity. Importantly, we capture the ways in which exploration and exploitation are not only paradoxical and contradictory forces in the knowledge creation process, but can also serve as substitutes and complements (Gupta et al. 2006). First, the model captures the opposing variance-inducing and variance-reducing characteristics of exploration and exploitation, respectively. Moreover, we also model the contrasting lagged realization of the payoffs from investing in exploration, versus the instantaneous benefits of investing in

exploitation. Second, we consider that exploration and exploitation act as substitutes, as both can increase the mean innovation performance outcome. Third, we recognize that exploration and exploitation may also be considered as complements, which provide synergies in the knowledge creation process. We incorporate this dynamic interdependence between exploration and exploitation through the inclusion of the effect of absorptive capacity. This feature of our model allows us to capture how prior investments in exploration and exploitation impact the effectiveness of future investments in exploration and exploitation.

Our results provide insights to aid managers identify effective strategies for pursuing knowledge creation and innovation under various risk objectives. While the analytical literature largely assumes a risk-averse, failure-sensitive decision-maker, the process of innovation and knowledge creation must necessarily assume that managers are likely to be risk-seeking and failure-tolerant in many cases. We examine the impact of these varying risk objectives on optimal temporal ambidexterity strategies. In a key result of the paper, we demonstrate that adopting a policy to “fail-early, fail-cheap” (Thomke 2001) may in-fact be the optimal strategy for a risk-averse manager, given a small initial pool of innovation solutions. Beyond establishing the boundary conditions for this sequential explore-exploit strategy, in which exploration precedes exploitation (Rothaermel and Deeds 2004), we are also able provide analytic insights on the conditions under which alternate dynamic knowledge creation strategies may endogenously arise. For example, we provide the conditions under which, contrary to commonly held wisdom to continuously reduce uncertainty, it may be optimal to employ a sequential strategy in which exploitation precedes exploration, such that the uncertainty

of innovation outcomes continues to evolve towards the end of the planning horizon (Leonardi 2011). Finally, by incorporating the lagged effects of exploration, and modeling the decision maker's dynamic risk preferences and short-term versus long-term objectives (Das and Teng 1997, Fried and Slowik 2004), we provide insights for managers faced with inter-temporal risk-taking decisions. Importantly, we demonstrate that the lagged effects exploration may act as either an incentive or disincentive to invest, depending on whether the manager is risk-averse versus risk-seeking, respectively (Abdellaoui et. al 2011). Collectively, our analytical results enable us to provide insights on the optimal sequencing of knowledge creation activities in the innovation process. As a managerial aid, our analytical results also highlight the optimal starting and/or stopping points which should be observed in order to avoid succumbing to possible exploration or exploitation learning traps (Levitt and March 1988, Levinthal and March 1993).

While our normative approach provides substantial insights, future opportunities exist to extend the model analytically. First, we recognize that other modeling features could be investigated. For example, there is an opportunity to examine the manager's budget creation and funding control mechanisms (Chao et al. 2009). Second, we do not consider the effects of market risk or exogenous uncertainty. As a contribution to the literature on ambidexterity, we believe our model provides a basis for developing testable empirical hypotheses of the four proposed temporal ambidexterity strategies: fixed exploration-dominant, fixed exploitation-dominant, sequential explore-exploit and sequential exploit-explore. Future empirical work which builds on our proposed dynamic knowledge creation model is needed to further develop an understanding of temporal ambidexterity. For example, a longitudinal study across one or many innovation teams

could provide a basis for testing the conditions under which different temporal ambidexterity patterns are observed (Van de Ven and Polley 1992).

## **CHAPTER 4**

### **A DIFFERENTIAL GAME MODEL OF EXPLORATION AND EXPLOITATION UNDER CO-OPETITION**

#### **4.1 Introduction**

A technology-based firm achieves a competitive advantage by developing an innovation which demonstrates technologically superior product performance relative to the technology alternatives currently available in the market (Ali et al. 1993, Sorescu et al. 2003, Sood and Tellis 2009). For example, in the disk drive industry firms compete on technical performance as measured by relative disk gate density (Lerner 1997, Franco et al. 2010) or relative processing speed (Khanna 1995). Since a firm's ability to innovate is largely a function of its knowledge resources and capabilities, competitive advantage is derived by a firm's ability to effectively and efficiently generate, acquire and deploy knowledge. However, for firms which operate in high-tech, high-risk, capital-intensive domains, often knowledge-sharing, resource-sharing, and risk-sharing are important strategies to consider (Gnyawali and Park 2011). Furthermore, a firm may even need to consider the option of knowledge-sharing with its direct competitors (Hamel 1991, Khanna et al. 1998, Baum et al. 2000, Spencer 2003). A classic example is the joint development of the LCD technology by Sony Corporation and Samsung Electronics, two firms that are direct competitors in the flat screen display market (Gnyawali and Park 2011). The strategy of collaborating with competitors is referred to as co-opetition (Walley 2007). Research suggests that firms which participate in cooperative development and knowledge-sharing achieve greater levels of innovation (Cellini and

Lambertini 2002, Spencer 2003, McGill and Santoro 2009). However, the resource based view of the firm proposes that competitive advantage is sustained from having high levels of proprietary information. Therefore, if a firm elects to participate in a knowledge-sharing partnership with its competitor, a critical decision is the extent to which it should focus on its private benefits, which drives its relative competitive advantage (Barney 1991), versus the collective benefits of the cooperative partnership (Kanter 1994), both of which can contribute to improved firm performance. For the remainder of the paper, we introduce the term *co-opetitive partner* to refer to a firm which partners with one of its direct competitors.

In this paper, we introduce a differential game model to examine the incentives which drive a firm to invest in its own knowledge creation as well as knowledge-sharing with its co-opetitive partner. Our model of co-opetitive learning allows us to examine how both competition and cooperation impact the rate, timing and sequence of the knowledge creation activities which a firm undertakes. We provide insights on how a firm should optimally manage its knowledge sharing exchanges, in order to balance the potential benefits of cooperation, against the potential threat of being “out-learned and out-competed by the competitor-partner” (Gnyawali and Park 2011, p. 657). Our analytical results provide a better understanding of the motivating factors which drive empirically observed alliance dysfunctions, wherein organizations delay knowledge-sharing and withhold information from their alliance partners (Hamel 1991, Khanna et al. 1998, Müller 2010).

Research suggests that a firm may limit the scope of its knowledge-sharing exchanges, as one mechanism for protecting against excessive knowledge loss and the



potential loss of competitive advantage (Oxley and Samson 2004). We contribute to this dialogue by examining how a firm's participation in different types of knowledge sharing exchanges may affect the incentives to exchange knowledge with, or to withhold knowledge from, a co-opetitive partner. Specifically, we consider two types of knowledge-sharing alliances: (i) an exploration alliance, which is focused on risk-taking and discovering new technologies, and (ii) an exploitation alliance, which is focused on risk-reduction and on refining and extending existing technologies (Colombo et al. 2006, Im and Rai 2008). We extend our earlier single firm mean-variance model of exploration and exploitation learning, described in Chapter 3, and introduce a differential game model to examine the moderating impact of competition and cooperation on a firm's optimal knowledge creation strategy. Rothaermel and Deeds (2004) show that alliances which undertake a specific sequence of activities, initially focusing on exploration and then later focusing on exploitation, outperform those which do not follow this sequence. However, they do not consider the possibility that the alliance partners are also direct competitors. We contribute to the research streams on exploration/exploitation, learning alliances, and co-opetition by examining the optimal sequencing of exploration and exploitation activities undertaken by co-opetitive partners, and provide examples where the typical explore-then-exploit sequential strategy, versus other alternate sequential strategies, may be optimal. We show how the optimal sequence of knowledge creation from exploration and exploitation and a firm's incentives and disincentives for knowledge sharing are influenced by: (i) the firm's participation in an exploration versus an exploitation alliance, (ii) the competitive performance payoff regime, and (iii) the firm's knowledge creation capabilities and the initial knowledge with which it is

endowed. Collectively, our findings contribute to a better understanding of the challenge and opportunities for a firm faced with the strategic imperative of engaging in a knowledge-sharing partnership with its competitor. By extrapolating from our analytic results, we offer managerial guidelines to determine “*how much*” and *what* knowledge should be shared, *when*, *with whom*, and *under what conditions*” (Loebecke et al. 1999).

## **4.2 Related Literature**

### **4.2.1 Innovation and Competition**

Two competing theories of the incentives for innovation under competition are the Arrowian and Schumpeterian hypotheses. The Arrowian hypothesis states that competition encourages innovation while the Schumpeterian hypothesis proposes that monopoly market power encourages innovation. Cellini and Lambertini (2002) examine a dynamic model of research and development investment under competition. They find support for the Arrowian hypothesis that innovation efforts aimed at product differentiation increase with competition. Oraopoulos and Kavadias (2008) develop a two-period model to determine how competing firms choose between investing in an unexplored versus an explored technological domain. They find that competition motivates exploration and more intense resource allocation. On the other hand, Boudreau et al. (2011) empirically find that competition discourages innovation efforts. Research also examines the impact of a firm’s relative competitive standing as an incentive for innovation. Lerner (1997) finds that firms which do not have a competitive advantage have greater incentive to innovate. It is important to recognize that in Lerner (1997), Cellini and Lambertini (2002) and Oraopoulos and Kavadias (2008), the decision maker selects and/or optimizes the level of effort for a single mode of innovation. In contrast, in

our dynamic game, we consider how a firm's competitive environment impacts its incentive to balance and sequence its investments in both exploration and exploitation knowledge creation activities throughout the innovation process.

#### **4.2.2 Innovation and Cooperation**

Several studies also examine the benefits of knowledge sharing and cooperation to improve innovation outcomes (McGill and Santoro 2009, Hora and Dutta 2012). Deeds and Hill (1996) show that technical and commercial success is increasing in the level of alliance participation for biotechnology firms. Spencer (2003) finds that, in the flat panel display industry, innovation performance improves with the degree of knowledge that a firm shares. Research therefore recognizes that a critical incentive for investing in knowledge creation is the availability of external sources of technical knowledge, referred to as the "technological opportunity" (Kamien and Schwartz 1975, Cohen and Levinthal 1990). Consistent with this view, Cellini and Lambertini (2002) develop an analytical model and demonstrate that research efforts aimed at product differentiation are increasing in R&D cooperation. Rothaermel and Alexandre (2009) propose that firm's innovative performance is a function of a balance between internal and external knowledge sourcing, as well as a balance between exploration and exploitation.

#### **4.2.3 Innovation and Learning under Co-opetition**

While the benefits of R&D cooperation and alliance formation are well accepted, for innovation and product differentiation, research also exists that highlights the challenges and potential disadvantages of knowledge-sharing. The disincentives for knowledge-sharing are particularly pertinent in cases where a firm must cooperate with its direct competitor (Hamel 1991, Inkpen 2000). Research has begun to examine the

tradeoffs involved in balancing a firm's cooperative incentives versus its competitive disincentives for sharing knowledge. Ozkan (2009) introduces two models to examine the incentives for competitive versus cooperative knowledge development. In the cooperative Stackelberg game model, they consider a leader and a follower that jointly develop knowledge and also enter the market together as a joint venture, based on the cumulative knowledge generated by the partnership. In the competitive Stackelberg game model, they consider a one-way transfer of knowledge in which the source firm has the option to sell or license knowledge to a recipient firm. After the knowledge transfer, both firms pursue additional knowledge development separately and then compete in the marketplace. They find that the recipient firm's decision to purchase knowledge is not only a function of the price of the knowledge, but also depends on the customer's valuation of the knowledge and the probability of successful development. In contrast to Ozkan (2009), however, we consider two-way knowledge sharing. Furthermore, our model also differs in that we assume that each firm freely reveals its knowledge to its competitor, without receiving compensation. As a result of this modeling assumption, we are able to examine and gain a better understanding of a firm's incentives for free-revealing and participating in open source innovation with its competitor (Von Hippel and von Krogh 2003, 2006). The two-way knowledge-sharing feature of our model allows us to capture the positive benefits, and also the negative consequences of engaging in knowledge-sharing partnerships. As such, we explain how a firm balances the tradeoff between the synergistic benefits knowledge-sharing, which can improve the donor firm's innovative performance, against the potential negative impacts of knowledge-sharing, if

the recipient is able to use the knowledge gained, to the donor firm's eventual competitive disadvantage (Khanna et al. 1998, Loebecke et al. 1999).

To date the literature on co-opetition has been largely focused on conceptual analysis and empirical testing. However, Loebecke et al. (1999) propose a game theoretic modeling framework may be developed to enable researchers to gain additional insights on the issue of co-opetition. Furthermore, in their call for future research Loebecke et al. (1999) also suggest the importance of including "time" as an additional aspect to be studied. Therefore, we contribute to the literature on co-opetition by introducing a differential game model of exploration and exploitation, which allows us to analytically examine the incentives for a firm to dynamically invest in knowledge creation and knowledge-sharing with its competitor. The differential game model allows us to consider the dynamic implications of a firm's initial competitive advantage, as well as the fact that, while a firm may find it necessary to share knowledge, the optimal timing for selectively revealing and sharing this knowledge may vary in different scenarios. Therefore, our dynamic results contribute to the open questions not only of "*how much*" and "*what knowledge should be shared*" but, importantly, we also provide insights on "*when*" and "*under what conditions*" knowledge-sharing with a co-opetitive partner is beneficial (Loebecke et al. 1999).

#### **4.2.4 Exploration, Exploitation and Risk-Taking in Innovation**

Our model differs from the related analytical knowledge-sharing models in the literature in that we model two different types of knowledge-sharing activities, while also considering the impact of competition. Firstly, we consider a case in which firms participate in an exploitation-based knowledge-sharing alliance which provides the

benefit of improved reduction of technical uncertainty (Koza and Lewin 2000, Colombo 2006). Exploitation activities include testing and refinement of existing knowledge, which allows a firm to reduce the level of technical uncertainty. However, exploitation tends to produce more incremental innovations (March 1991, Fleming 2001, He and Wong 2004). Secondly, we also consider a case in which the competing firms participate in an exploration-based knowledge-sharing alliance, which provides both firms with knowledge related to the discovery of new technologies (Koza and Lewin 2000, Rothaermel and Deeds 2004, Colombo 2006). Exploration is related to the generation of variation and trial-and-error learning, to discover new and radical innovations, which can possibly achieve quantum leaps in technical performance. However, by employing more experimental technologies, the ability to reliably predict the outcome of the innovation is lower (March 1991, Bohn 1994, He and Wong 2004).

In their call for future research Loebecke et al. (1999) stress the importance of considering the incentives which drive participation in co-opetitive partnerships. Therefore, to incorporate risk-taking incentives, we build on March's (1991) framework of competition for relative position in a right-tail race to be the best, versus in a left-tail race to avoid finishing last. Several other studies have examined the impact of performance incentives on risk-taking choices. Cabral (2003) examines the choice between a low variance versus a mean preserving high variance innovation. He shows that when firms compete on relative performance, in an infinite period race, the laggard selects the risky technology. Tsetlin et al. (2004) also show that, to increase the likelihood of winning, a competitor should maximize variability if it is in a weak relative position, which they define as having low mean performance. On the other hand, Boyle

and Shapira (2012) analyze a player's competitive strategies in the Jeopardy Tournament of Champions game and find that a leader, not the laggard, is more inclined to select the riskier strategy. While the above literature examines the impact of competition and a firm's competitive advantage on the propensity for pursuing risk-taking activities, such as exploration, versus pursuing risk-reducing activities, such as exploitation, it does not consider the impact of collaboration. In contrast, we consider how a firm's competitive performance regime influences its optimal knowledge-sharing strategy. Collectively, we contribute to the research on organizational learning, alliances, and co-opetition by examining the optimal sequencing of a firm's pursuit of exploration, exploitation and knowledge-sharing within a co-opetitive partnership.

#### **4.3 A Model of Exploration and Exploitation with Cooperation and Competition**

##### **4.3.1 A Dynamic Model of Knowledge Creation**

We consider two technology-based firms, each with a manager responsible for generating the knowledge necessary to develop a novel innovation. Knowledge creation occurs over a fixed time horizon  $t \in [0, T]$ , where 0 is the initial time of the innovation development phase. The terminal time  $T$  when this phase of development concludes is given and is the same for both firms. At the end of the development phase,  $t=T$ , both firms compete against each other in the market on the basis of the technical performance of the innovations each has developed.

For convenience, in the remainder of the paper we consider two firms,  $j=1,2$ , where we refer to  $j=1$  as the focal firm, and  $j=2$  as the rival firm. The manager of each firm can invest in both exploration-based and exploitation-based modes of knowledge creation. Let the control variables  $i(t)$ ,  $e(t)$  and  $I(t)$ ,  $E(t)$  denote the rates of investment in

exploration and exploitation at time  $t \in [0, T]$ , for the focal (small letters) and the rival firm (capital letters), respectively. All control variables are bounded below by zero. Operating costs are incurred as each firm invests in exploration and exploitation. The parameters reflecting the costs of exploration and exploitation are defined as  $c_0, c_1$  and  $C_0, C_1$ , for the focal and the rival firm, respectively. We make no assumptions about the comparative cost of exploration versus exploitation a-priori. Consistent with the literature, we assume quadratic cost functions for both modes of knowledge creation to reflect the diseconomies of scale due to the disruption and coordination of larger-scale knowledge creation activities at any single instant in time.

Given the uncertain outcome of the innovation development process for both firms, we model the technical performance of each innovation, which is achievable at time  $t$  as a random variable,  $v(t)$  and  $V(t)$ , for the focal and the rival firm, respectively (March 1991, Carrillo and Gaimon 2004, Oraopoulos and Kavadias 2008). The performance outcomes are assumed to be distributed according to two separate normal distributions,  $v(t) \sim N(\mu(t), \sigma(t))$  and  $V(t) \sim N(M(t), S(t))$ , for the focal firm and the rival firm, respectively. The variables  $\mu(t)$ ,  $\sigma(t)$  and  $M(t)$ ,  $S(t)$  denote the mean and variance of the performance of the innovation developed by each firm. These variables represent the expected performance of the innovation as well as the uncertainty associated with realizing that performance at time  $t$ , respectively. At the initial time,  $\mu(0) > 0$  and  $M(0) > 0$  represent the mean technical performance of each firm's innovation at the beginning of the development phase, which is driven by each firm's prior knowledge and experience. The variance  $\sigma(t)$  and  $S(t)$  reflects the range and predictability of the possible performance outcome achievable by the innovation developed by each firm at time  $t$ . At



the initial time,  $\sigma(0)>0$  and  $S(0)>0$  represent the initial variance in the technical performance of the innovation for each firm. The initial variance is high if the previous ideas pursued prior to the current development phase were highly novel, and low if they are based on existing and well understood technology.

For each firm, investments in knowledge creation dynamically alter the mean and variance of the probability distribution that characterizes the technical performance of the innovation (March 1991). Investing in exploitation generates improvements in the mean technical performance. Moreover, exploitation refines the firm's existing knowledge and resolves uncertainty by focusing on routinizing processes, generating repeatable outcomes and identifying unpromising alternatives. Therefore, exploitation is said to have a variance-reducing effect on the innovation process. In contrast, exploration of new knowledge allows for the discovery of novel technological approaches which can substantially increase the upside potential of the technical performance of the innovation. However, by employing more experimental technologies, the ability to predict the outcome of the innovation process is lower, so that exploration is said to have a variance-increasing effect on the innovation process (March 1991, Bohn 1994). The variance-reducing impact of exploitation and the variance-increasing impact of exploration are well accepted (March 1991, Fleming 2001, He and Wong 2004).

The mathematical relationships captured in Equations (1) and (2) reflect the impact of exploration and exploitation on the mean and variance of the technical performance of the innovation at time  $t$ , for firm  $j$ , where  $j = 1$ . Let  $dG/dg$  denote the first order derivative of  $G$  with respect to  $g$ . Note that  $\alpha_0, \beta_0, \alpha_1, \beta_1$  are all positive constants. The coefficients  $\alpha_0>0$  and  $\beta_0>0$  in Equation (1) represent the extent to which a unit of

exploration versus a unit of exploitation, at time  $t$ , positively impacts the level of the mean technical performance of the innovation. Consistent with our empirical findings from Chapter 2, in the second term of Equation (2), we model the variance-reducing impact of exploitation at time  $t$  (March 1991, He and Wong 2004), as a reduction in the variance of the technical performance of the innovation at that time. On the other hand, exploration at time  $t$  is assumed to have a variance-increasing effect at that time (March 1991, He and Wong 2004), as reflected in the first term in Equation (2). The coefficients  $\alpha_1 > 0$  and  $\beta_1 > 0$  in Equation (2) represent the extent to which a unit of exploration versus a unit of exploitation impacts the variance of the technical performance of the innovation for the focal firm  $j=1$ . The analogous equations hold for firm  $j=2$  with  $a_0, b_0, a_1, b_1$ , as positive constants.

#### **4.3.2 Two-way Knowledge-Sharing of Exploration and Exploitation Knowledge**

A key incentive for investing in knowledge creation is the availability of external technical knowledge (Kamien and Schwartz 1975, Cohen and Levinthal 1990, Khanna et al. 1998, Baum et al. 2000, Spencer 2003). Equations (1) and (2) also capture the mechanism by which the two competing firms cooperate by participating in a two-way knowledge-sharing agreement. Two elements determine the extent to which a firm benefits from its alliance partner's knowledge: (i) the cumulative knowledge of the source firm, and (ii) the degree to which the source firm commits to participating in the knowledge-sharing alliance. The degree of knowledge sharing participation is higher if a firm reveals more knowledge. The extent of knowledge-sharing by each firm reflects an up-front agreement between the two firms and is assumed given (i.e., determined prior to the initial time).

We recognize that a firm can participate in either an exploration or an exploitation-based alliance (Rothaermel and Deeds 2004, Colombo 2006, Lavie and Rosenkopf 2006, Hoang and Rothaermel 2010). Consistent with the empirical literature, our model focuses on the variance-enhancing versus the variance-reducing benefits which a firm derives from participating in an exploration versus an exploitation knowledge-sharing partnership, respectively (Schulz 2001, Rothaermel and Deeds 2004, Im and Rai 2008). Im and Rai (2008) empirically demonstrate the variance reducing benefits of an exploitation knowledge-sharing alliance. In Equation (2) below, we model the benefits of participating in an exploitation-based alliance, which reflects the benefits of knowledge-sharing between firms related to low-risk, short-term improvements and the refinement of existing systems (Im and Rai 2008). Specifically, the rate at which the focal firm reduces its own variance at time  $t$  is driven by three factors: (i) the focal (recipient) firm's rate of variance reduction due to its own exploitation efforts,  $e(t)$ , (ii) the rival (donor) firm's cumulative variance-reduction efforts, as reflected by a smaller variance at time  $t$ ,  $S(t)$ , and (iii) the rival (donor) firm's level of participation in the exploitation knowledge-sharing alliance,  $Y$ . Reinganum (1981) employs a similar knowledge sharing parameter in a model of innovation and with rivalry. We assume diminishing returns to the knowledge sharing benefit such that  $Y \in (0,1)$ . Importantly, as the rival (donor) firm's variance,  $S(t)$ , is smaller (less uncertainty) then the marginal reduction in the focal (recipient) firm's variance with respect to  $e(t)$  (given by  $\beta_1 S(t)^{-Y} e(t)$  as it appears in the second term on Equation (2)), improves (note the minus sign in the power of  $S(t)$ ). This modeling assumption is consistent with empirical research, which finds that as a donor organization's level of codification of existing knowledge increases (i.e. as donor firm's

performance variability is reduced (smaller  $S(t)$ ), then the outflows of codified exploitation-related knowledge from the donor is higher (Schulz 2001). The analogous equations hold for the rival firm, where the focal firm's level of participation in the exploitation alliance is given by  $\gamma$ .

Im and Rai (2008) also empirically examine the variance enhancing benefits of an exploration knowledge-sharing alliance. In Equation (2) below we model the benefits of participating in an exploration-based alliance, which reflects the benefits of knowledge-sharing between firms related to experimentation and innovation discovery involving significant risk and uncertainty. Specifically, the rate at which the focal firm can increase its own variance at time  $t$  is driven by three factors: (i) the rate of variation generation due to the focal firm's own exploration efforts,  $i(t)$ , (ii) the rival (donor) firm's cumulative variance enhancement efforts, as reflected by a larger variance at time  $t$ ,  $S(t)$ , and (iii) the rival (donor) firm's level of participation in the exploration knowledge-sharing alliance,  $X$ . We assume diminishing returns to the knowledge sharing benefit such that  $X \in (0,1)$ . As the pool of alternative solutions generated by the donor firm increases ( $S(t)$  gets larger), then the focal (recipient) firm's ability to increase its own variance from exploration, as given by the first term in Equation (2), increases (note the power of  $S(t)$  is positive). This modeling assumption is consistent with empirical research which finds that, as the variability and uniqueness of the donor's knowledge is higher (i.e. as donor firm's performance variability is enhanced (larger  $S(t)$ ), then the outflows of exploration-related knowledge from the donor are higher (Schulz 2001). The analogous equations hold for the rival firm, where the focal firm's level of participation in the exploration alliance is given by  $\psi$ .

Reflecting the above discussion, we obtain Equations (1) and (2), below. From Equation (1), given  $\mu(0)$ ,  $i(t)$ , and  $e(t) \geq 0$ , the non-negativity of  $\mu(t)$  is satisfied. However, from Equation (2) since it is possible that the non-negativity of  $\sigma(t)$  could be violated we introduce a non-negativity constraint.

$$\frac{d\mu(t)}{dt} = \alpha_0 i(t) + \beta_0 e(t) \quad (1)$$

$$\frac{d\sigma(t)}{dt} = \alpha_1 i(t) S(t)^x - \beta_1 e(t) S(t)^{-y} \quad (2)$$

Beyond enhancing the knowledge creation abilities of the recipient firm, a key benefit of knowledge sharing is the opportunity for cost sharing. Therefore, we assume that a firm does not incur any additional costs for benefitting from the partner's knowledge or for participating in the knowledge-sharing alliance (Von Hippel and Von Krogh 2006). That is, we assume that the firm, which initially generates the knowledge, freely reveals this knowledge and does not derive any profit from knowledge-sharing. However, each firm derives a non-financial benefit because it reciprocally improves its marginal effectiveness of knowledge creation as a result of the shared knowledge. Lastly, note that a firm does not derive any benefit from the knowledge-sharing alliance unless it independently invests in its own internal exploration and exploitation efforts (Hoang and Rothaermel 2010), which is necessary to integrate and deploy the shared knowledge.

### 4.3.3 Competing on Knowledge in a Right-tail and Left-tail Race

The payoff which the focal firm realizes at the end of the planning horizon is a function of the relative terminal performance of the innovations developed by both firms,  $f(v(T), V(T))$ . Empirical research confirms the impact of the relative technical advantage on firm performance, with respect to increased revenues, increased firm market value or

an increased likelihood of firm survival (Ali 1994, Sorescu et al. 2003, Franco et al. 2010). For example, Franco et al. (2010) measure relative performance advantage and relative competitive position in the disk drive industry, based on the disk density of a focal firm relative to the highest disk density available in the market within the same year. They show that the likelihood of firm survival in the subsequent year is positively related to the relative technical advantage.

Leveraging the aspiration theory (March and Shapira 1992, Wiseman and Bromiley 1996), we model two different competitive performance regimes. Based on March's (1991) framework of competition for relative position in a right-tail race versus in a left-tail race, a firm is either rewarded for having the best relative performance or penalized for having the worst relative performance, respectively. This is also consistent with the risk-taking literature which refers to survival versus aspiration targets, respectively (March and Shapiro 1992, March 1988). Highlighting the managerial biases reflected by these performance regimes, Dickson and Giglierano (1986) refer to a manager's concern for either "sinking the boat" by performing below a certain minimum target, or "missing the boat" by failing to achieve high performance outcomes. Roels and Su (2013) model a similar set of performance constructs, which they refer to as behind-averse versus ahead-seeking behaviors, respectively.

In the first competitive performance regime, each firm considers the innovation's upside potential and strives to demonstrate the highest performance relative to its competitor. Given the two terminal performance distributions, the likelihood of the focal firm achieving the highest relative competitive position increases in relation to the mean and variance of its own performance (March 1991, Singh and Fleming 2010), but

decreases in relation to the mean and variance of the rival firm's performance (March 1991). Therefore, we define the expected terminal payoff for a focal firm that seeks to achieve the best technical performance of its innovation relative to its competitor, as given in Equation (3) below. Our model assumption is consistent with the results demonstrated by Tsetlin et al. (2004) who, using simulation, confirm that higher mean and variance are substitutes in increasing the likelihood of winning.

In the second competitive performance regime, a firm is penalized for having the lowest relative performance and therefore tries to minimize this potential downside risk. Given the two terminal performance distributions, the focal firm's competitive advantage increases in relation to the mean and decreases in relation to the variance of its own performance, but decreases in relation to the mean, and increases in relation to the variance of the rival firm's performance (March 1991). Therefore, we define the expected terminal payoff for the focal firm that seeks to avoid the worst relative performance, as given in Equation (4), below. This notion of managerial risk-avoidance is consistent with the threat-rigidity literature (Staw et al. 1981) which suggests that a firm pursues activities which reduce risk and uncertainty in order to limit the probability of poor outcomes and maximize the likelihood of firm survival, or to avoid finishing last (March 1991).

To summarize, the focal firm maximizes the objective as defined in Equation (5) subject to the dynamics in Equations (1) and (2), and the non-negativity constraints on  $i(t)$ ,  $e(t)$  and  $\sigma(t)$ . Analogous expressions hold for the rival firm ( $j=2$ ). That is, the manager of each firm optimally invests in exploration and exploitation in order to: (i) minimize the cumulative expenditures incurred for knowledge creation over the innovation process and (ii) maximize the expected terminal payoff. The objective

function is captured in Equation (5) with  $f[v(T), V(T)]$  for each performance regime as given in in Equations (3) ((4)) under competition to achieve the best relative performance (avoid the worst relative performance). We solve the model separately for each performance regime represented by Equations (3) and (4).

$$f_1[v(T), V(T)] = \mu(T) + \sigma(T) - M(T) - S(T) \quad (3)$$

$$f_2[v(T), V(T)] = \mu(T) - \sigma(T) - M(T) + S(T) \quad (4)$$

$$-\frac{1}{2} \int_0^T \{c_0 i^2(t) + c_1 e^2(t)\} dt + f[v(T), V(T)] \quad (5)$$

We model knowledge sharing under competition as an open loop differential game (Sethi and Thompson 2000). Under the open loop solution each firm pre-commits to an optimal knowledge creation strategy, in the sense that they do not modify their actions once the development process begins (Gaimon 1989). This reflects a strategy in which each firm assesses its internal capability and external environment to determine an optimal plan, which it commits to and then executes over the development horizon. It has been shown under certain conditions, that a player in a differential game can improve its performance by playing an open loop game (Jorgensen 1982, Gaimon 1989)<sup>1</sup>. Koza and Lewin (2000) also highlight the importance of strategic intent during the alliance formation process, as critical to success.

#### 4.4 Optimal Solutions

The model is solved using optimal control theory (Sethi and Thompson 2000, Hartl and Sethi 1984). A summary of notation appears in Table C1 of Appendix C. We

---

<sup>1</sup> Jorgensen (1982) cites instances where both firms playing open-loop strategies outperform both playing a closed loop strategy including Starr and Ho (1969), Mukundan and Elsner (1975), Case (1979). Gaimon (1989) finds that both firms playing an open loop strategy dominates both playing the closed loop approach. However, if one firm plays an open game against a closed loop competitor then Gaimon shows the open loop player wins.



introduce four adjoint variables for each firm. For the focal firm,  $j=1$ ,  $\lambda_\mu(t)$ ,  $\lambda_\sigma(t)$ ,  $\lambda_M(t)$ , and  $\lambda_S(t)$  represent the marginal value to its objective of a unit increase in its own mean and variance at time  $t$ , as well as a unit change in the rival's mean and variance at time  $t$ , respectively. The analogous variables for the rival firm,  $j=2$ , are  $L_\mu(t)$ ,  $L_\sigma(t)$ ,  $L_M(t)$ , and  $L_S(t)$ . For the focal firm,  $j=1$ , let  $\eta_\sigma(t)$  ( $\eta_S(t)$ ) represent the Lagrange multiplier associated with the non-negativity constraint on  $\sigma(t)$  ( $S(t)$ ). The required optimality conditions and complementary slackness conditions for  $\eta_\sigma(t)$  and  $\eta_S(t)$  are provided in Appendix C. However, for simplicity, in the remainder of the paper, we assume that the constraints are non-binding, so that  $\sigma(t)>0$ ,  $S(t)>0$ ,  $\eta_\sigma(t)=0$  and  $\eta_S(t)=0$  hold for  $t \in [0, T]^2$ . Throughout the remainder of the paper "\*" refers to an optimal solution. The optimality conditions and proofs appear in the Appendix C. In the remainder of this section, we provide analytical insights on the factors which lead to various optimal dynamic knowledge creation and knowledge-sharing strategies. In many cases, we are able to derive the key insights analytically. In other instances, we rely on numerical analysis. Details of the numerical analyses are provided in Appendix C.

To simplify the analysis which follows, we consider the optimal solutions for four special cases (see Table 4.1). In Case 1, both firms participate in an exploration-only alliance in which they exchange knowledge related to variance generation, that is  $\psi>0$ ,  $X>0$ ,  $\gamma=0$ ,  $Y=0$ . Furthermore, in Case 1A both firms compete to achieve the best performance, based on the terminal payoff structure as given in Equation (3). In contrast, in Case 1B both firms compete to avoid the worst performance, based on the terminal payoff structure as given in Equation (4). In Case 2, both firms participate in an

---

<sup>2</sup> For reasonable parameter settings, throughout the extensive numerical analysis conducted these constraints were not violated.

exploitation-only alliance in which they exchange knowledge related to variance reduction, that is  $\psi=0$ ,  $X=0$ ,  $\gamma>0$  and  $Y>0$ . In Case 2A both firms compete to achieve the best performance, based on the terminal payoff structure as given in Equation (3), while in Case 2B both firms compete to avoid the worst performance based on the terminal payoff structure as given in Equation (4).

**Table 4.1: Special Cases**

Case	Knowledge Sharing Alliance Type	Competitive Performance Regime
1A	Exploration-only Knowledge-sharing alliance $\psi>0$ , $X>0$ , $\gamma=0$ , $Y=0$	Compete to achieve the Best performance, $f_1$
1B		Compete to avoid the Worst performance, $f_2$
2A	Exploitation-only Knowledge-sharing alliance $\psi=0$ , $X=0$ , $\gamma>0$ , $Y>0$ .	Compete to achieve the Best performance, $f_1$
2B		Compete to avoid the Worst performance, $f_2$

In each of the four cases (1A, 1B, 2A, 2B) we consider the optimal rates at which the manager should pursue both exploration and exploitation throughout the planning horizon. We define the knowledge creation strategy as *front-loaded*, when the rate of the knowledge creation activity is positive and decreasing over all time ( $de^*(t)/dt < 0$  or  $di^*(t)/dt < 0$  for  $t \in [0, T]$ ), so that the peak rate occurs at the initial time (Thomke and Fujimoto 2000, Thomke 2001, Carrillo and Gaimon 2004, Ozkan 2009, Xiao 2012). In contrast, we define the knowledge creation strategy as *back-loaded* when the rate of the knowledge creation is positive and increasing over all time ( $de^*(t)/dt > 0$  or  $di^*(t)/dt > 0$  for  $t \in [0, T]$ ), so that the peak rate occurs at the terminal time (Carrillo and Gaimon 2004, Fixson and Marion 2012, Xiao 2012). In addition, we define a *delay* strategy as one in which the maximum rate of knowledge creation optimally occurs at some time during the planning horizon (inverse U-shaped). Finally, we introduce the notion of a knowledge creation strategy referred to as a *bookend* strategy, whereby the manager initially front-

loads and also later back-loads its investment in knowledge creation, (i.e., the rate of knowledge creation optimally reaches a minimum at some time during the planning horizon and is U-shaped).

#### 4.4.1 Case 1: Knowledge Sharing in an Exploration-only Alliance

In Theorem 1 below, we introduce the optimal solutions for the rates of exploration and exploitation as well as the marginal values of the variance for the focal firm ( $j=1$ ) within an exploration-only alliance. Recall that in Case 1, the firms exchange knowledge as they pursue their individual variance generation activities, so that  $\psi > 0$ ,  $X > 0$ ,  $\gamma = 0$ ,  $Y = 0$  (see Table 1).

**THEOREM 1.** *Within an exploration-only alliance, the optimal rates of exploration and exploitation and the corresponding rates of change at time  $t$  satisfy Equations (6) and (7). The rates of change of the marginal values of an additional unit of the focal firm's and its rival's mean and variance at time  $t$  satisfy Equations (8) and (9).*

$$e^*(t) = \text{Max} \left\{ \frac{\lambda_\mu(t)\beta_0 - \lambda_\sigma(t)\beta_1}{c_1}, 0 \right\}; \frac{de(t)^*}{dt} = \frac{-\frac{d\lambda_\sigma(t)}{dt}\beta_1}{c_1} \quad (6)$$

$$i^*(t) = \text{Max} \left\{ \frac{\lambda_\mu(t)\alpha_0 + \lambda_\sigma(t)\alpha_1 S(t)^X}{c_0}, 0 \right\}; \frac{di(t)^*}{dt} = \frac{\frac{d\lambda_\sigma(t)}{dt}\alpha_1 S(t)^X + \lambda_\sigma(t)\alpha_1 X \frac{dS(t)}{dt} S(t)^{X-1}}{c_1} \quad (7)$$

$$\frac{d\lambda_\mu(t)}{dt} = 0; \frac{d\lambda_\sigma(t)}{dt} = -a_1 I(t) \psi \sigma(t)^{\psi-1} \lambda_s(t) \quad (8)$$

$$\frac{d\lambda_M(t)}{dt} = 0; \frac{d\lambda_s(t)}{dt} = -\alpha_1 i(t) X S(t)^{X-1} \lambda_\sigma(t) \quad (9)$$

The optimal rates of exploration and exploitation are driven by the marginal value to the focal firm of its own mean and variance ( $\lambda_\mu(t)$ ,  $\lambda_\sigma(t)$ ). Furthermore, the marginal value to the focal firm of its own variance ( $\lambda_\sigma(t)$ ) is in turn a function of the marginal

value to the focal firm of an increase in its rival's variance ( $\lambda_S(t)$ ). In addition, the optimal rates of exploration ( $i^*(t)$ ) and exploitation ( $e^*(t)$ ) depend on the focal firm's marginal effectiveness of exploration and exploitation ( $\alpha_0, \beta_0, \alpha_1, \beta_1$ ) and the knowledge shared by the rival firm ( $S(t)^X$ ). To further examine the optimal solutions for an exploration-only alliance, as given in Theorem 1, we consider subcases 1A and 1B (see Table 4.1). Corollary 1A corresponds to Case 1A, whereas, Corollary 1B corresponds to Case 1B.

**COROLLARY 1A.** *For a firm operating under competition to be the best performer within an exploration-only alliance the marginal value to the focal firm of its own mean is positive and of the rival firm's mean is negative, while the marginal value of the firm's own variance and of the rival firm's variance satisfies one of the following three cases for  $t \in [0, T]$ : (I)  $\lambda_\sigma(t) > 0$ ,  $d\lambda_\sigma(t)/dt \geq 0$ ,  $\lambda_S(t) < 0$ ,  $d\lambda_S(t)/dt \leq 0$  for  $t \in [0, T]$ ; (II)  $\lambda_\sigma(t) \leq 0$ ,  $d\lambda_\sigma(t)/dt \geq 0$ ,  $\lambda_S(t) < 0$ ,  $d\lambda_S(t)/dt \geq 0$  for  $t \in [0, t_1]$  and  $\lambda_\sigma(t) > 0$ ,  $d\lambda_\sigma(t)/dt \geq 0$ ,  $\lambda_S(t) < 0$ ,  $d\lambda_S(t)/dt < 0$  for  $t \in (t_1, T]$  where  $t_1 \in (0, T)$  or (III)  $\lambda_\sigma(t) > 0$ ,  $d\lambda_\sigma(t)/dt \leq 0$ ,  $\lambda_S(t) \geq 0$ ,  $d\lambda_S(t)/dt \leq 0$  for  $t \in [0, t_2]$  and  $\lambda_\sigma(t) > 0$ ,  $d\lambda_\sigma(t)/dt > 0$ ,  $\lambda_S(t) < 0$ ,  $d\lambda_S(t)/dt \leq 0$  for  $t \in (t_2, T]$  where  $t_2 \in (0, T)$ . The optimal solutions for  $\lambda_\sigma(t)$  and  $\lambda_S(t)$  for  $t \in [0, T]$  are illustrated in Figure C1 (Appendix C).*

From the proof of Corollary 1A, the terminal marginal value to the focal firm of its own variance,  $\lambda_\sigma(T)$ , is positive reflecting the benefits of increasing the likelihood of an extreme right tail outcome, thereby gaining the competitive advantage of being the best performer. Conversely, the terminal marginal value to the focal firm of an increase in the rival firm's variance,  $\lambda_S(T)$ , is negative since a reduction in the rival firm's variance reduces the likelihood that the rival firm achieves an extreme right tail outcome. However, as the focal firm invests in increasing its own variance,  $\sigma(t)$ , it also contributes

to improving the variance of its co-opetitive partner. Therefore, as described below, the firm must manage the various tensions which arise as a consequence of investing in its own knowledge creation, and potentially ceding its competitive advantage due to the knowledge gained by its co-opetitive partner.

In Case 1A-I, under competition to be the best within an exploration alliance, the marginal value to the focal firm of increasing its own variance,  $\lambda_\sigma(t)$ , is positive. However, the focal firm recognizes that its rival can benefit from the outflow of its own knowledge,  $(-a_1 I(t) \psi \sigma(t)^{\psi-1} \lambda_s(t)$  in Equation(8)). Therefore, although it is valuable for the focal firm to increase its variance ( $\lambda_\sigma(t) > 0$ ), it is optimal to delay that increase until later in the planning horizon ( $d\lambda_\sigma(t)/dt > 0$ ). The focal firm postpones increasing its own variance as this also improves the rival firm's variance, to the focal firm's competitive detriment. This apprehension about sharing knowledge with its competitor gives rise to the focal firm adopting a dysfunctional alliance behavior, in which it delays and limits the extent of knowledge creation and knowledge-sharing for purposes of competitive deterrence (Hamel 1991, Khanna et al. 1998, Müller 2010). On the other hand, while the marginal value to the focal firm of an increase in its rival's variance is negative ( $\lambda_s(t) < 0$ ), the focal firm also recognizes that it can benefit from the knowledge inflows from its co-opetitive partner  $(-\alpha_1 i(t) X S(t)^{X-1} \lambda_\sigma(t)$  in Equation (9)). Therefore, although the marginal value of the rival firm's variance is negative, the focal firm benefits from any knowledge inflows from the co-opetitive partnership earlier in the planning horizon ( $d\lambda_s(t)/dt \leq 0$ ). It is important to recognize that as the terminal time approaches when the firms compete in the marketplace, the focal firm's incentives with respect to improving its own variance ( $\lambda_\sigma(t) > 0$ ,  $d\lambda_\sigma(t)/dt \geq 0$ ) versus its incentives for improving its rival's variance ( $\lambda_s(t) < 0$ ,

$d\lambda_s(t)/dt \leq 0$ ) are not only opposing, but also becoming less aligned. Analytically, the diverging magnitude and direction of the marginal value functions seem to capture Hamel's (1991) observation that within strategic alliances, "as the firm moved nearer and nearer its goal of independence, it successively raised the 'price' for its continued participation in the alliance" (p. 88).

Case 1A-II demonstrates that even in an exploration-only alliance, a focal firm competing to be the best performer may have an incentive to decrease its own variance early in the planning horizon ( $\lambda_\sigma(t) < 0$ ,  $t \in [0, t_1]$ ). Mathematically, the solution given in Case 1A-II applies when either of the following hold: (i) the marginal effectiveness of the rival firm's variance generation capabilities or the rate of exploration of the rival increases (i.e.,  $a_1$  or  $I(t)$  is larger) or (ii) the cumulative variance or the degree of knowledge-sharing of the focal firm increases (i.e.,  $\sigma(t)$  or  $\psi$  is larger). Therefore, this case occurs when the potential competitive loss resulting from the knowledge outflow from the focal firm is large ( $-a_1 I(t) \psi \sigma(t)^{\psi-1} \lambda_s(t)$ ).

Finally, under Case 1A-III, the marginal value to the focal firm of increasing its co-opetitive partner's variance,  $\lambda_s(t)$ , is first positive and decreasing, and then later non-positive and decreasing. This is an interesting result as it demonstrates a possible scenario under which the focal firm initially has an incentive to enhance its rival's variance in a right-tail race ( $\lambda_s(t) > 0$  for  $t \in [0, t_2]$ ). Furthermore, when the focal firm has an incentive to increase its rival's variance, it is also optimal to increase its own variance earlier in the planning horizon ( $d\lambda_\sigma(t)/dt < 0$ ). Consequently, under these conditions the focal firm's incentives for delaying knowledge creation and knowledge-sharing are reversed relative to Case 1A-I. Moreover, the focal firm's incentives with respect to improving its own

variance ( $\lambda_\sigma(t) > 0$ ) versus its incentives for improving its rival's variance ( $\lambda_S(t) > 0$ ) are aligned, since both are initially positive and non-increasing early in the planning horizon ( $\lambda_\sigma(t) > 0, d\lambda_\sigma(t)/dt \leq 0, \lambda_S(t) > 0, d\lambda_S(t)/dt \leq 0$  for  $t \in [0, t_2]$ ). Mathematically, the solution given in Case 1A-III applies when either of the following hold: (i) the marginal effectiveness of the focal firm's variance generation capabilities or its rate of exploration increases (i.e.,  $\alpha_1$  or  $i(t)$  is larger) or (ii) the cumulative variance or the degree of knowledge-sharing of the rival firm increases (i.e.,  $S(t)$  or  $X$  is larger). Therefore, this case occurs when the benefits gained from the knowledge inflow from the rival firm are large ( $-\alpha_1 i(t) X S(t)^{X-1} \lambda_\sigma(t)$ ). This suggests that when the focal firm holds the weaker competitive position in a right-tail race, its incentive to participate in the exploration knowledge-sharing alliance increases.

Based on Theorem 1, and since the marginal value of the focal firm's own mean is positive and constant, the optimal rates of both exploration and exploitation follow from the solutions for the marginal value of the variance,  $\lambda_\sigma(t)$ , in Corollary 1A. For Cases 1A-I and III, since the marginal value of the variance is positive, ( $\lambda_\sigma(t) > 0$  for  $t \in [0, T]$ ), it is always optimal to invest in exploration, ( $i^*(t) > 0$  for  $t \in [0, T]$ ). However, for Case 1A-II, although the firm is competing to be the best performer, during the initial periods for which the marginal value of the variance,  $\lambda_\sigma(t)$ , is negative, if the firm's rate of variance generation,  $\alpha_1$ , outpaces the mean improvements from exploration,  $\alpha_0$ , then it may not be optimal to explore, ( $i^*(t) = 0$ ). Therefore, there may exist an optimal starting time, which we define as  $t_{\text{start}}$ , before which the firm does not participate in the exploration knowledge-sharing alliance.

The rate of exploration by the focal firm,  $i^*(t)$ , also depends on the rate of increase in the rival's variance,  $dS(t)/dt$ . Suppose the marginal value to the focal firm of its own variance is positive. If the rival's variance is increasing, ( $dS/dt > 0$ ), it is better for the focal firm to invest in exploration at an increasing rate in order to wait until the rival firm has undertaken more exploration and has a larger variance so that the focal firm can benefit from a larger knowledge pool later in the horizon. From Equation (7), if the rival's variance is increasing, ( $dS/dt > 0$ ), given the solution for the marginal value of the variance,  $\lambda_\sigma(t)$ , in Corollary 1A, it follows that the focal firm's investment in exploration is optimally back-loaded (see Figure 4.1). The optimal rate of the focal firm's exploration is also back-loaded in Case 1A-II. However, for an initial interval of time in which the marginal value of the variance is negative, ( $\lambda_\sigma(t) < 0$ ,  $t \in [0, t_1)$ ), exploration may not optimally occur if the variance effect dominates the marginal improvement in the mean ( $\alpha_0 \lambda_\mu(t) + \alpha_1 \lambda_\sigma(t) S(t)^X < 0$ ,  $i^*(t) = 0$ ), (see Figure 4.2). Finally, in Case 1A-III, as the potential knowledge inflows from the rival increase ( $-\alpha_1 i(t) X S(t)^{X-1} \lambda_\sigma(t)$ ), exploration may follow a bookend strategy, in which the manager initially front-loads and later back-loads the rate of exploration, attaining its minimum value at some time during the planning horizon (U-shaped), as illustrated in Figure 4.3. Therefore, the manager optimally pursues a *Back-load Explore-Front-load Exploit* strategy in Cases 1A-I and 1A-II, or a *Bookend Explore-Delay Exploit* strategy in Case 1A-III. The solutions are illustrated in Figures 4.1, 4.2 and 4.3, respectively. For Cases 1A-I, II and III, note that when the marginal value of the variance is positive, ( $\lambda_\sigma(t) > 0$ ), if the focal firm's rate of variance reduction,  $\beta_1$ , outpaces the mean improvements from exploitation,  $\beta_0$ , it is



possible that exploitation may not optimally occur under competition to be the best, if  $\lambda_\mu(t)\beta_0 - \lambda_\sigma(t)\beta_0 < 0$  holds ( $e^*(t)=0$ ).




Figure 4.1: Case 1A-I





Figure 4.2: Case 1A-II



**Figure 4.3: Case 1A-III**

**COROLLARY 1B.** *For a firm operating under competition to avoid being the worst performer within an exploration-only alliance, the marginal value to the focal firm of its own mean is positive and of the competitor's mean is negative, while the marginal value of the firm's own variance and of the competitor's variance satisfies one of the following three cases for  $t \in [0, T]$ : (I)  $\lambda_\sigma(t) < 0$ ,  $d\lambda_\sigma(t)/dt \leq 0$ ,  $\lambda_S(t) > 0$ ,  $d\lambda_S(t)/dt \geq 0$  for  $t \in [0, T]$ ; (II)  $\lambda_\sigma(t) < 0$ ,  $d\lambda_\sigma(t)/dt \geq 0$ ,  $\lambda_S(t) \leq 0$ ,  $d\lambda_S(t)/dt \geq 0$  for  $t \in [0, t_1]$  and  $\lambda_\sigma(t) < 0$ ,  $d\lambda_\sigma(t)/dt < 0$ ,  $\lambda_S(t) > 0$ ,  $d\lambda_S(t)/dt > 0$  for  $t \in (t_1, T]$ ; (III)  $\lambda_\sigma(t) \geq 0$ ,  $d\lambda_\sigma(t)/dt \leq 0$ ,  $\lambda_S(t) > 0$ ,  $d\lambda_S(t)/dt \leq 0$  for  $t \in [0, t_2]$  and  $\lambda_\sigma(t) < 0$ ,  $d\lambda_\sigma(t)/dt \leq 0$ ,  $\lambda_S(t) > 0$ ,  $d\lambda_S(t)/dt > 0$  for  $t \in (t_2, T]$ . The optimal solutions for  $\lambda_\sigma(t)$  and  $\lambda_S(t)$  for  $t \in [0, T]$  are illustrated in Figure C3(Appendix C).*

Under Case 1B-I, in a left-tail race, the marginal value to the focal firm of increasing its co-opetitive partner's variance,  $\lambda_S(t)$ , is positive and increasing. The focal firm is motivated to increase its rival's variance since this increases the likelihood that the rival realizes the worst performance. On the other hand, the marginal value to the focal firm of increasing its own variance,  $\lambda_\sigma(t)$ , is negative and decreasing. That is, while it is valuable for the focal firm to reduce its variance ( $\lambda_\sigma(t) < 0$ ), it is more valuable to reduce the variance later in the planning horizon. Therefore, as in Corollary 1A, the focal firm's

incentives with respect to changes in its own variance versus changes in the rival firm's variance become increasingly divergent over time (Hamel 1991).

Under Case 1B-II, the marginal value to the focal firm of the rival's variance,  $\lambda_S(t)$ , is initially negative. This is an interesting result since it provides possible conditions under which the incentive for the focal firm to reduce its own variance, as well as its incentive to reduce the variance of the rival firm, are initially aligned (i.e., both are negative and increasing ( $\lambda_\sigma(t) < 0$ ,  $d\lambda_\sigma(t)/dt \geq 0$ ,  $\lambda_S(t) < 0$ ,  $d\lambda_S(t)/dt \geq 0$ ,  $t \in [0, t_1)$ ). The solution given in Case 1B-II applies when either of the following hold: (i) the marginal effectiveness of the focal firm's variance generation capabilities or its rate of exploration increases (i.e.,  $\alpha_1$  or  $i^*(t)$  is larger) or (ii) as either the cumulative variance or the degree of knowledge-sharing from the rival firm increases (i.e.  $S(t)$  or  $X$  is larger). When the rival's variance is large in an exploration knowledge-sharing alliance this poses a competitive hazard to the focal firm as this also increases its own variance, which reduces its likelihood of avoiding an extreme left tail outcome. Therefore, there is an initial incentive for the focal firm to reduce its own variance and to reduce the variance of the rival firm ( $\lambda_\sigma(t) < 0$ ,  $d\lambda_\sigma(t)/dt \geq 0$ ,  $\lambda_S(t) \leq 0$ ,  $d\lambda_S(t)/dt \geq 0$  for  $t \in [0, t_1)$ ). However, later in the planning horizon the incentives for the focal firm to improve its own knowledge versus improving its rival's knowledge are again divergent ( $\lambda_\sigma(t) < 0$ ,  $d\lambda_\sigma(t)/dt < 0$ ,  $\lambda_S(t) > 0$ ,  $d\lambda_S(t)/dt > 0$  for  $t \in (t_1, T]$ ). Importantly, in Cases 1B-I and 1B-II, since the marginal value to the focal firm of increasing its variance is negative, ( $\lambda_\sigma(t) < 0$ ,  $t \in [0, T]$ ), the focal firm may optimally select not to participate in the exploration knowledge-sharing alliance if the variance effect dominates the marginal improvement in the mean ( $\alpha_0\lambda_\mu(t) + \alpha_1\lambda_\sigma(t)S(t)^X < 0$ ,  $i^*(t) = 0$ ).

Finally, Case 1B-III demonstrates that even when a focal firm competes to avoid being the worst performer it may have an incentive to increase its own variance, as well as an incentive to increase its rival's variance, early in the planning horizon ( $\lambda_{\sigma}(t) > 0$  for  $t \in [0, t_2]$ ). The solution given in Case 1B-III applies when either of the following hold: (i) the marginal effectiveness of the rival's variance generation capabilities or its rate of exploration increases (i.e.  $a_1$  or  $I^*(t)$  is larger) or (ii) the cumulative variance or the degree of knowledge-sharing of the focal firm increases (i.e.,  $\sigma(t)$  or  $\psi$  is larger). This suggests that when the focal firm holds the weaker competitive position in a left-tail race it has a greater incentive to participate in the exploration knowledge-sharing alliance. Under these conditions, the focal firm may deliberately increase its variance even further, and worsen its competitive position, in order to limit to the rival's competitive advantage.


The following optimal solutions for the rates of exploration and exploitation follow analytically from Theorem 1 and Corollary 1B-I. Given that the marginal value to the focal firm of increasing its own variance,  $\lambda_{\sigma}(t)$ , is negative and decreasing, exploitation optimally follows a backload strategy. Since both firms are competing to avoid the worst performance, it is reasonable to assume that each firm's variance is optimally decreasing, ( $d\sigma/dt < 0$ ,  $dS/dt < 0$ ). Since the marginal value to the focal firm is negative, under comparable marginal cost and knowledge creation effectiveness for both exploration and exploitation, since exploration (exploitation) is increasing (decreasing) in  $\lambda_{\sigma}(t)$ , it is reasonable to assume that the  $dS/dt$  term is negative, and dominates the expression in Equation (7). Therefore, exploration follows a backload strategy. Intuitively, the focal firm seeks to avoid being the worst performer, and therefore the firm has an incentive to delay its own exploration, and therefore defers its participation in the

exploration alliance,  $i(t)$ , until later in the planning horizon when the co-opetitive partner has resolved some of its initial uncertainty,  $S(t)$ . Therefore, under the conditions described above, the manager optimally pursues a *Backload Explore-Backload Exploit* strategy in Cases 1B-I, as illustrated in Figure 4.4 below (see Appendix C, Table C2, Figure C4).


We numerically analyze and highlight several interesting optimal solutions that arise from Cases 1B-II and III, as contrasted with Case 1B-I (see Appendix C, Table C2, Figures C5). Suppose the focal firm,  $j=1$ , has access to very cheap and effective exploration capabilities so that its rate of exploration increases (i.e., cost,  $c_0$ , is small, and mean exploration improvement,  $\alpha_0$  is large, so that  $i^*(t)$  is larger), then Case 1B-II holds. From Case 1B-II, for the focal firm,  $j=1$ , the marginal value of its own variance,  $\lambda_\sigma(t)$ , is initially negative and increasing, and then later negative and decreasing and the marginal value of an increase in the rival firm's variance,  $\lambda_S(t)$ , is initially negative and increasing, and then later positive and increasing. However, for the rival firm,  $j=2$ , as the variance generation capabilities or the rate of exploration of firm  $j=1$  increases ( $i^*(t)$  is larger), the solution for firm  $j=2$  follows Case 1B-III.

From Case 1B-III, for the rival firm  $j=2$  the marginal value of its own variance,  $L_S(t)$ , is initially positive and decreasing, and then later negative and decreasing, and firm  $j=2$  optimally front-loads exploration, ( $I^*(t)$ ), and backloads exploitation, ( $E^*(t)$ ). However, for firm  $j=2$ , later in the planning horizon, when the marginal value of its own variance,  $L_S(t)$ , is negative and decreasing, it may not be optimal for firm  $j=2$  to explore ( $I^*(t)=0$ ) (see Figure 4.5). From Equation (8), when firm  $j=2$  does not invest in exploration,  $I^*(t)=0$ , the dynamic incentives for the focal firm,  $j=1$ , is eliminated, and so


that the marginal value to firm  $j=1$  of its own variance is constant,  $d\lambda_\sigma/dt=0$ , when this occurs later in the planning horizon. Also recall that the marginal value of the focal firm's own variance,  $\lambda_\sigma(t)$ , is initially negative and increasing, and then later negative and decreasing then, so that for the focal firm, exploitation optimally follows a bookend strategy. Furthermore, given that the marginal value to firm  $j=1$  of its own variance is constant,  $d\lambda_\sigma/dt=0$ , later in the planning horizon, the rate of exploitation is also constant later in the planning horizon. Clearly, since the marginal value of firm  $j=1$ 's own variance is negative for the entire planning horizon, ( $\lambda_\sigma(t)<0$  for  $t\in[0,T]$ ), then firm  $j=1$  is driven not to invest in exploration. However, given the conditions assumed in this scenario, that is, since the marginal effectiveness of the focal firm's capabilities to improve the mean from exploration are large (i.e.,  $\alpha_0$  is large,  $\alpha_0\lambda_\mu(t)+\alpha_1\lambda_\sigma(t)S(t)^X>0$ ), then exploration optimally occurs, ( $i^*(t)>0$ ). Furthermore, given that the marginal value to firm  $j=1$  of its own variance is constant later in the planning horizon,  $d\lambda_\sigma/dt=0$ , then the  $dS/dt$  term, which is negative later in the horizon, dominates so that firm  $j=1$  pursues exploration at an increasing rate later in the planning horizon, as illustrated in Figure 4.6.



**Figure 4.4: Case 1B-I for firm  $j=1$**




**Front-load Exploration**




**Back-load Exploitation**

**Figure 4.5: Case 1B-III for firm  $j=2$**



**Bookend Exploration**



**Bookend Exploitation**

**Figure 4.6: Case 1B-II for Firm  $j=1$**

#### 4.4.2 Case 2: Knowledge Sharing in an Exploitation-only Alliance

In Theorem 2 below, we introduce the solutions for the focal firm's optimal rates of exploration and exploitation, within an exploitation-based knowledge sharing alliance. Recall that in Case 2, both firms exchange knowledge as they pursue their individual variance reduction activities, so that  $\psi=0$ ,  $X=0$ ,  $\gamma>0$ ,  $Y>0$ .

**THEOREM 2** *Within an exploitation-only knowledge sharing alliance, the optimal rates and the rates of change of exploration and exploitation at time  $t$ , satisfy Equations (10)*

and (11). The rates of change of the marginal values of an additional unit of the focal firm's and its rival's mean and variance at time  $t$  satisfy Equations (12) and (13).

$$e^*(t) = \text{Max} \left\{ \frac{\lambda_\mu(t)\beta_0 - \lambda_\sigma(t)\beta_1 S(t)^{-Y}}{c_1}, 0 \right\}; \frac{de(t)^*}{dt} = \frac{-\frac{d\lambda_\sigma(t)}{dt}\beta_1 S(t)^{-Y} + \lambda_\sigma(t)\beta_1 Y \frac{dS(t)}{dt} S(t)^{-Y-1}}{c_1} \quad (10)$$

$$i^*(t) = \text{Max} \left\{ \frac{\lambda_\mu(t)\alpha_0 + \lambda_\sigma(t)\alpha_1}{c_0}, 0 \right\}; \frac{di^*(t)}{dt} = \frac{\frac{d\lambda_\sigma(t)}{dt}\alpha_1}{c_1} \quad (11)$$

$$\frac{d\lambda_\sigma(t)}{dt} = -b_1 E(t) \gamma \sigma(t)^{-\gamma-1} \lambda_s(t) \quad (12)$$

$$\frac{d\lambda_s(t)}{dt} = -\beta_1 e Y S(t)^{-Y-1} \lambda_\sigma(t) \quad (13)$$

To examine the optimal solutions for an exploitation-only alliance, as given in Theorem 2, we consider subcases 2A and 2B, as shown in Table 4.1. Corollary 2A corresponds to Case 2A in which both firms compete to achieve the best performance within an exploitation-only alliance, while Corollary 2B corresponds to Case 2B in which both firms compete to avoid the worst performance, within the exploitation-only alliance.

**COROLLARY 2A.** *For a firm operating under competition to be the best performer within an exploitation-only alliance, the marginal value to the focal firm of its own mean is positive and of the competitor's mean is negative, while the marginal value of the firm's own variance and of the rival firm's variance satisfies one of the following three cases for  $t \in [0, T]$ : (I)  $\lambda_\sigma(t) > 0$ ,  $d\lambda_\sigma(t)/dt \geq 0$ ,  $\lambda_s(t) < 0$ ,  $d\lambda_s(t)/dt \leq 0$  for  $t \in [0, T]$ ; (II)  $\lambda_\sigma(t) \leq 0$ ,  $d\lambda_\sigma(t)/dt \geq 0$ ,  $\lambda_s(t) < 0$ ,  $d\lambda_s(t)/dt \geq 0$  for  $t \in [0, t_1]$  and  $\lambda_\sigma(t) > 0$ ,  $d\lambda_\sigma(t)/dt \geq 0$ ,  $\lambda_s(t) < 0$ ,  $d\lambda_s(t)/dt < 0$  for  $t \in (t_1, T]$  where  $t_1 \in (0, T)$  or (III)  $\lambda_\sigma(t) > 0$ ,  $d\lambda_\sigma(t)/dt \leq 0$ ,  $\lambda_s(t) \geq 0$ ,  $d\lambda_s(t)/dt \leq 0$  for  $t \in [0, t_2]$  and  $\lambda_\sigma(t) > 0$ ,  $d\lambda_\sigma(t)/dt > 0$ ,  $\lambda_s(t) < 0$ ,  $d\lambda_s(t)/dt \leq 0$  for  $t \in (t_2, T]$  where  $t_2 \in (0, T)$ .*



**COROLLARY 2B.** *For a firm operating under competition to avoid being the worst performer within an exploitation-only alliance, the marginal value to the focal firm of its own mean is positive and of the competitor's mean is negative, while the marginal value of the firm's own variance and of the competitor's variance satisfies one of the following three cases for  $t \in [0, T]$ : (I)  $\lambda_\sigma(t) < 0$ ,  $d\lambda_\sigma(t)/dt \leq 0$ ,  $\lambda_S(t) > 0$ ,  $d\lambda_S(t)/dt \geq 0$  for  $t \in [0, T]$ ; (II)  $\lambda_\sigma(t) < 0$ ,  $d\lambda_\sigma(t)/dt \geq 0$ ,  $\lambda_S(t) \leq 0$ ,  $d\lambda_S(t)/dt \geq 0$  for  $t \in [0, t_1]$  and  $\lambda_\sigma(t) < 0$ ,  $d\lambda_\sigma(t)/dt < 0$ ,  $\lambda_S(t) > 0$ ,  $d\lambda_S(t)/dt > 0$  for  $t \in (t_1, T]$ ; (III)  $\lambda_\sigma(t) \geq 0$ ,  $d\lambda_\sigma(t)/dt \leq 0$ ,  $\lambda_S(t) > 0$ ,  $d\lambda_S(t)/dt \leq 0$  for  $t \in [0, t_2]$  and  $\lambda_\sigma(t) < 0$ ,  $d\lambda_\sigma(t)/dt \leq 0$ ,  $\lambda_S(t) > 0$ ,  $d\lambda_S(t)/dt > 0$  for  $t \in (t_2, T]$ .*

From Corollary 2A, under competition to be the best performer within an exploitation alliance, exploitation is most likely to optimally occur under Case 2A-II, for the initial interval of time in which the marginal value of the variance is negative, ( $\lambda_\sigma(t) < 0$ ,  $t \in [0, t_1]$ ), so that  $\beta_0 \lambda_\mu(t) - \beta_1 \lambda_\sigma(t) S(t)^{-Y} > 0$  holds. Therefore, in a right-tail competition, the focal firm is motivated to participate in the exploitation knowledge-sharing alliance when the knowledge outflow from the focal firm is larger ( $-b_1 I(t) \gamma \sigma(t)^{-\gamma-1} \lambda_S(t)$ ), which is increasing as the focal firm's variance is smaller ( $\sigma(t)^{-\gamma-1}$ ). This suggests that when the focal firm holds the weaker competitive position in the right-tail race, it may deliberately engage in the exploitation alliance in order to decrease its own variance even further, and forgoing its competitive position, in order to worsen to the rival's competitive position.

From Corollary 2B, under competition to avoid being the worst performer within an exploitation alliance, exploitation is least likely to occur under Case 2B-III, for the initial interval of time in which the marginal value of the variance is positive, ( $\lambda_\sigma(t) > 0$ ,  $t \in [0, t_1]$ ), so that  $\beta_0 \lambda_\mu(t) - \beta_1 \lambda_\sigma(t) S(t)^{-Y} < 0$  holds. Therefore, in a left-tail competition, the

focal firm is not motivated to participate in the exploitation knowledge-sharing alliance when the knowledge inflow from the rival firm is larger  $(-\beta_1 i(t) Y S(t)^{-Y-1} \lambda_o(t))$ , which is increasing as the rival firm's variance is smaller  $(S(t)^{-Y-1})$ . This suggests that when the focal firm holds the weaker competitive position in a left-tail race, this decreases its incentive to participate in the exploitation knowledge-sharing alliance.

#### 4.5 Managerial Insights and Conclusions

In this essay, we extend the model of temporal ambidexterity, introduced in Essay 1, to include considerations of the strategic imperatives of competition and cooperation. In particular, we introduce a dynamic optimization model of knowledge-sharing between two rival firms. A firm can participate in knowledge-sharing in order to explore new technological opportunities, as well as to improve the ability to exploit its existing capabilities. As such, we consider two alternative types of alliances: (i) an exploration knowledge-sharing alliance and (ii) an exploitation knowledge-sharing alliance. Based on the analysis of our model we are able to examine how a firm's collaborative and competitive objectives, as well as its innovation capabilities, impact its optimal knowledge creation strategy. Empirically, Yang et al. (2010) demonstrate that source firms benefit from their own knowledge contributions when a "spillover knowledge pool" is formed from which they can benefit. However, through our stylized model, we illustrate the challenges of leveraging the benefits of knowledge spillovers within a co-opetitive partnership. Specifically, we demonstrate the tension which co-opetitive partners face as they try to manage their knowledge creation and knowledge sharing activities in order to "maximize incoming while minimizing outgoing knowledge spillovers" (Alexy et al. 2013) and to ensure that their competitor's do not "outlearn"

them. This apprehension about sharing knowledge with a potential competitor is aptly captured by Hamel in one manager's comment, from a case study of strategic alliances: "Whatever they learn from us, they'll use against us worldwide" (Hamel, 1991, p.87). Our results demonstrate how such concerns within knowledge-sharing alliances, can give rise to firms adopting dysfunctional alliance behaviors in which they may delay and limit the extent of knowledge sharing, for purposes of competitive deterrence (Hamel 1991, Khanna et al. 1998, Müller 2010).

Our results show that a firm's dysfunctional alliance behavior is increasing in cases in which the losses from the knowledge outflow from the focal firm dominates the benefits it realizes from the knowledge inflow from its co-opetitive partner, (i.e., if the firm believes it has more to lose than it has to gain). This suggests that competitive leaders are more likely to adopt dysfunctional alliance behaviors, while laggards are more likely to participate in the alliance. Consistent with this insight, in a study of biotechnology firms, Shan (1990) empirically confirms that laggards are more likely than leaders to participate in cooperative arrangements.

Hamel's (1991) observes that within strategic alliances, "as the firm moved nearer and nearer its goal of independence, it successively raised the 'price' for its continued participation in the alliance" (p. 88). In support of this view, we find that in several instances a firm's incentives for improving its own performance versus improving its co-opetive partner's performance, are not only opposing, but may also become increasingly divergent as the firm approaches the time to compete in the market.

Our results suggest the boundary conditions, with respect to both the alliance structure and competitive regime, under which the typical "explore-exploit" strategy may

be optimal within an exploration knowledge-sharing alliance (Rothaermel and Deeds 2004). Furthermore, our results may provide additional insights to support these empirical findings, by analytically demonstrating the conditions, with respect to competitive regime, relative knowledge stocks and relative learning capabilities, under which this solution arises endogenously. Moreover, our remaining analytical solutions also provide the conditions under which alternate sequential strategies arise endogenously. Future research may empirically test and validate the conditions and the corresponding optimal solutions obtained.

## APPENDIX A

**Table A1: Range of Exploration-Exploitation Strategies**

% of Patents	High Exploration	Low Exploration
High Exploitation	23% (3072 of 13464)	33% (4385 of 13464)
Low Exploitation	15% (1962 of 13464)	30% (4045 of 13464)

**Table A2: Descriptive Statistics and Correlations**

**Full Sample: 13464 patents**

Variable	Mean	S.D	Min	Max	1	2	3	4	5
1 Failure	0.10	0.30	0	1					
2 Success	0.03	0.17	0	1	-0.06				
3 Explore	0.38	0.25	0	0.91	0.003	0.04			
4 Exploit	11.02	7.72	0	97	0.09	-0.05	0.09		
5 Prior Failures (in Five Years)	1.01	2.39	0	17	0.04	-0.02	0.02	-0.03	
6 Prior Successes (in Five Years)	1.52	3.53	0	22	-0.03	0.11	0.08	-0.03	0.34

**Sub-sample of publicly listed firms: 2836 patents**

Variable	Mean	S.D	Min	Max	1	2	3	4	5	6	7
1 Failure	0.08	0.27	0	1							
2 Success	0.04	0.19	0	1	-0.06						
3 Explore	0.41	0.27	0	0.89	0.01	0.03					
4 Exploit	10.92	7.04	1	68	0.13	-0.05	0.11				
5 Prior Failures (in 5 Years)	2.33	3.58	0	17	0.09	-0.03	0.06	0.03			
6 Prior Successes (in 5 Years)	5.19	5.8	0	22	-0.03	0.12	0.11	-0.01	0.24		
7 Assets	6456	11100	1.66	168259	0.01	0.04	-0.01	-0.02	-0.03	-0.05	
8 R&D Intensity	0.81	5.5	0.01	77	-0.02	-0.01	-0.05	-0.05	-0.07	-0.12	-0.01

**Table A3: Complementary Log-Log Model for Probability of Success and Failure**

	Probability of Success					Probability of Failure				
	(S1)	(S2)	(S3)	(S4)	(S5)	(F1)	(F2)	(F3)	(F4)	(F5)
Explore	0.34*** (0.08)	0.35*** (0.08)			0.34*** (0.07)	-0.04+ (0.02)	-0.07** (0.02)			-0.06* (0.03)
Exploit	-0.49*** (0.05)	-0.49*** (0.05)			-0.51*** (0.06)	0.33*** (0.03)	0.35*** (0.03)			0.35*** (0.03)
Explore X Exploit		0.01 (0.05)			0.02 (0.04)		0.11*** (0.02)			0.10*** (0.02)
Prior Success Experience			0.59*** (0.05)		0.61*** (0.06)			-0.19*** (0.04)		-0.19*** (0.04)
Prior Failure Experience			-0.54*** (0.09)		-0.53*** (0.09)			0.10** (0.03)		0.11** (0.03)
Exploit X Prior Failure				0.24*** (0.06)	-0.01 (0.05)				0.03 (0.03)	0.01 (0.03)
Explore X Prior Failure				-0.16+ (0.08)	-0.02 (0.08)				0.01 (0.02)	0.01 (0.02)
Exploit X Prior Success				-0.23*** (0.04)	0.09** (0.03)				-0.03 (0.04)	-0.00 (0.04)
Explore X Prior Success				0.14 (0.10)	-0.09* (0.04)				0.00 (0.02)	0.02 (0.02)
Constant	-4.77*** (0.34)	-4.77*** (0.34)	-4.85*** (0.34)	-4.91*** (0.35)	-4.75*** (0.34)	-3.17*** (0.15)	-3.18*** (0.15)	-3.03*** (0.15)	-2.95*** (0.15)	-3.23*** (0.15)
$N$	13385	13385	13464	13385	13385	13385	13385	13464	13385	13385
$\chi^2$	200	198	281	141	471	588	601	520	502	642
Prob > $\chi^2$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Log Likelihood	-1830	-1830	-1789	-1866	-1729	-3955	-3948	-4041	-4017	-3931

Standard errors in parentheses. +  $p < 0.10$ , \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ .  
Clustered by Assignee. Year and Technology Class Fixed Effects Included

**Table A4: Linear Regression for Mean and Variance of Citations Received**

	Variance					Mean				
	Absolute Deviation of Number of Citations Received					Number of Citations Received				
	(V1)	(V2)	(V3)	(V4)	(V5)	(M1)	(M2)	(M3)	(M4)	(M5)
Explore	0.04*** (0.01)	0.04*** (0.01)			0.03** (0.01)	0.06*** (0.01)	0.06*** (0.01)			0.05*** (0.01)
Exploit	-0.05*** (0.01)	-0.05*** (0.01)			-0.05*** (0.01)	-0.14*** (0.01)	-0.15*** (0.01)			-0.15*** (0.01)
Explore X Exploit		-0.01* (0.00)			-0.01* (0.00)		-0.04*** (0.01)			-0.03*** (0.00)
Prior Success Experience			0.13*** (0.02)		0.12*** (0.02)			0.22*** (0.02)		0.21*** (0.02)
Prior Failure Experience			-0.07*** (0.01)		-0.07*** (0.01)			-0.14*** (0.02)		-0.14*** (0.02)
Exploit X Prior Failure				0.03* (0.01)	0.01 (0.01)				0.05* (0.02)	0.01 (0.01)
Explore X Prior Failure				-0.02+ (0.01)	-0.02+ (0.01)				-0.03 (0.02)	-0.03* (0.01)
Exploit X Prior Success				-0.04* (0.01)	-0.02+ (0.01)				-0.05* (0.02)	-0.02 (0.01)
Explore X Prior Success				0.02 (0.02)	0.01 (0.01)				0.03 (0.02)	0.01 (0.02)
Constant	0.57*** (0.02)	0.57*** (0.02)	0.59*** (0.02)	0.54*** (0.02)	0.61*** (0.02)	-0.22*** (0.03)	-0.22*** (0.03)	-0.22*** (0.03)	-0.29*** (0.03)	-0.16*** (0.03)
N	13385	13385	13464	13385	13385	13385	13385	13464	13385	13385
R <sup>2</sup>	0.022	0.023	0.040	0.020	0.047	0.059	0.061	0.081	0.043	0.103
F	7.056	6.609	10.02	5.531	8.045	24.72	22.74	32.76	16.32	27.04

Standard errors in parentheses. +  $p < 0.10$ , \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ .

Clustered by Assignee. Year and Technology Class Fixed Effects Included

**Table A5: Quantile Regression for Citations Received**

	(1) 10 <sup>th</sup> Percentile	(2) 50 <sup>th</sup> Percentile	(3) 90 <sup>th</sup> Percentile
Explore	0.01 (0.02)	0.17*** (0.04)	1.12*** (0.17)
Exploit	-0.26*** (0.02)	-0.89*** (0.04)	-2.44*** (0.18)
Explore X Exploit	-0.03 (0.02)	-0.23*** (0.04)	-0.68*** (0.17)
Prior Success Experience	0.22*** (0.02)	1.13*** (0.05)	4.09*** (0.20)
Prior Failure Experience	-0.13*** (0.02)	-0.59*** (0.05)	-1.96*** (0.20)
Exploit X Prior Failure	0.07** (0.02)	0.19*** (0.05)	0.25 (0.21)
Explore X Prior Failure	0.01 (0.02)	-0.11* (0.04)	-0.53** (0.20)
Exploit X Prior Success	-0.09*** (0.02)	-0.21*** (0.04)	-0.32 (0.20)
Explore X Prior Success	-0.04+ (0.02)	0.04 (0.04)	0.30 (0.19)
Constant	1.66*** (0.07)	7.61*** (0.16)	25.28*** (0.59)
N	13385	13385	13385
Pseudo R2	0.0364	0.0885	-4.9023

Standard errors in parentheses. +  $p < 0.10$ , \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ .

Clustered by Assignee. Year and Technology Class Fixed Effects Included



**Table A6: Seemingly Unrelated Regression (SUR) for Citations Received**

	(1) Explore	(2) Exploit	(3) Mean Number of Citations Received	(4) Explore	(5) Exploit	(6) Variance Number of Citations Received
Prior Failure Experience	-0.02* (-2.4)	-0.05*** (-6.29)	-1.41*** (-14.52)	-0.02* (-2.44)	-0.05*** (-6.29)	-0.07*** (-10.40)
Prior Success Experience	0.09*** (9.81)	-0.00 (-0.15)	2.23*** (22.73)	0.09*** (9.81)	-0.00 (-0.15)	0.12*** (16.60)
Explore			0.40*** (4.61)			0.03*** (5.34)
Exploit			-1.38*** (-15.02)			-0.05*** (-7.59)
Explore X Exploit			-0.34*** (-3.96)			-0.01* (-2.12)
Exploit X Prior Failure			0.24** (2.59)			0.01 (1.87)
Explore X Prior Failure			-0.29** (-3.13)			-0.02*** (-3.67)
Exploit X Prior Success			-0.15 (-1.65)			-0.02*** (-3.49)
Explore X Prior Success			0.01 (0.11)			0.01* (2.08)
Constant	-0.00 (-0.01)	-0.00 (-0.40)	11.39*** (36.45)	-0.00 (-0.01)	-0.00 (-0.40)	0.61*** (25.96)
R2	0.007	0.003	0.178	0.007	0.003	0.047
$\chi^2$	99	48	2901	99	48	673
Prob > $\chi^2$	0.000	0.000	0.000	0.000	0.000	0.000

Standard errors in parentheses. +  $p < 0.10$ , \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ .  
Clustered by Assignee. Year and Technology Class Fixed Effects Included

**Table A7: Complementary Log-Log Model for Probability of Success**

	Probability of Success		
	(S1)	(S2)	(S3)
	Top 1%	Top 2%	Top 3%
Explore	0.55*** (0.15)	0.37*** (0.09)	0.34*** (0.07)
Exploit	-0.46*** (0.09)	-0.55*** (0.07)	-0.51*** (0.06)
Explore X Exploit	0.04 (0.07)	0.02 (0.05)	0.02 (0.04)
Prior Success Experience	0.82*** (0.07)	0.68*** (0.07)	0.61*** (0.06)
Prior Failure Experience	-0.75*** (0.15)	-0.66*** (0.13)	-0.53*** (0.09)
Exploit X Prior Failure	0.13* (0.06)	0.00 (0.06)	-0.01 (0.05)
Explore X Prior Failure	-0.05 (0.11)	0.03 (0.10)	-0.02 (0.08)
Exploit X Prior Success	0.08+ (0.04)	0.08* (0.03)	0.09** (0.03)
Explore X Prior Success	-0.14* (0.06)	-0.12* (0.04)	-0.09* (0.04)
Constant	-6.08*** (0.57)	-5.02*** (0.36)	-4.75*** (0.34)
N	13385	13385	13385
$\chi^2$	390	416	471
Prob > $\chi^2$	0.000	0.000	0.000
Log Likelihood	-773	-1309	-1729

Standard errors in parentheses. +  $p < 0.10$ , \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ .  
Clustered by Assignee. Year and Technology Class Fixed Effects Included

**Table A8: Zero Inflated Negative Binomial Model for Citations Received**

	Number of Citations Received				
	(1)	(2)	(3)	(4)	(5)
Explore	0.08*** (0.01)	0.08*** (0.01)			0.06*** (0.01)
Exploit	-0.16*** (0.01)	-0.18*** (0.01)			-0.17*** (0.01)
Explore X Exploit		-0.04*** (0.01)			-0.03*** (0.01)
Prior Success Experience			0.22*** (0.02)		0.21*** (0.01)
Prior Failure Experience			-0.14*** (0.02)		-0.14*** (0.02)
Exploit X Prior Failure				0.05* (0.02)	0.00 (0.01)
Explore X Prior Failure				-0.03 (0.02)	-0.02 (0.01)
Exploit X Prior Success				-0.03+ (0.01)	0.01+ (0.01)
Explore X Prior Success				0.02 (0.02)	-0.00 (0.01)
Constant	2.18*** (0.04)	2.18*** (0.04)	2.16*** (0.04)	2.11*** (0.04)	2.23*** (0.04)
<i>Inflation Factor</i>					
Explore	1.25* (0.51)	1.10+ (0.63)	0.63* (0.31)	0.56+ (0.31)	1.12* (0.49)
Exploit	1.89*** (0.26)	1.96*** (0.26)	1.82*** (0.16)	1.87*** (0.16)	1.88*** (0.25)
Constant	-9.35*** (1.64)	-12.46 (35.03)	-7.21*** (0.46)	-7.37*** (0.50)	-9.16*** (1.51)
Inalpha Constant	-0.14*** (0.029)	-0.14*** (0.029)	-0.18*** (0.026)	-0.12*** (0.02)	-0.20*** (0.02)
N	13385	13385	13385	13385	13385
$\chi^2$	1489	1490	1223	1056	1698
Prob > $\chi^2$	0.000	0.000	0.000	0.000	0.000
Log Likelihood	-40631	-40619	-40508	-40790	-40318

**Table A9: Complementary Log-Log Model, Sub-sample of Publicly Listed Firms**

	Probability of Success		Probability of Failure	
	(S1)	(S2)	(F1)	(F2)
Explore	0.20 (0.21)	0.22 (0.19)	-0.14 (0.0949)	-0.13 (0.09)
Exploit	-0.59*** (0.13)	-0.57*** (0.13)	0.57*** (0.11)	0.56*** (0.11)
Explore X Exploit	-0.10 (0.12)	-0.07 (0.13)	0.24** (0.08)	0.25** (0.08)
Prior Success Experience	0.58*** (0.10)	0.62*** (0.10)	-0.15* (0.06)	-0.15* (0.06)
Prior Failure Experience	-0.31* (0.15)	-0.27+ (0.14)	0.05 (0.06)	0.05 (0.06)
Exploit X Prior Failure	-0.11+ (0.06)	-0.10+ (0.06)	0.004 (0.07)	0.00 (0.07)
Explore X Prior Failure	0.04 (0.13)	0.05 (0.13)	0.01 (0.04)	0.01 (0.04)
Exploit X Prior Success	0.12+ (0.06)	0.13* (0.05)	-0.11 (0.07)	-0.11 (0.06)
Explore X Prior Success	-0.06 (0.07)	-0.07 (0.07)	0.05 (0.04)	0.04 (0.04)
<b>Control Variables</b>				
Assets		0.17* (0.08)		0.00 (0.03)
R&D Intensity		0.32** (0.12)		0.00 (0.06)
Constant	-4.838*** (1.004)	-5.18*** (1.003)	-3.250*** (0.392)	-3.14*** (0.49)
N	2825	2782	2825	2782
$\chi^2$	287	510	476	397
Prob > $\chi^2$	0.000	0.000	0.000	0.000
Log Likelihood	-411	-400	-715	-696

Standard errors in parentheses. +  $p < 0.10$ , \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ .  
Clustered by Assignee. Year and Technology Class Fixed Effects Included

**Table A10: Complementary Log-Log Model for Probability of Success with Alternate Exploit Measure**

	(S1)	(S2)	(S3)	(S4)	(S5)
Explore	1.04*** (0.30)	1.04*** (0.30)			0.97*** (0.28)
Exploit II	0.06 (0.04)	0.07 (0.04)			0.04 (0.05)
Explore X Exploit II		0.05 (0.05)			0.02 (0.05)
Prior Success Experience			0.72*** (0.06)		0.73*** (0.07)
Prior Failure Experience			-0.82*** (0.15)		-0.90*** (0.16)
Exploit II X Prior Failure				-0.10* (0.05)	0.09 (0.05)
Explore X Prior Failure				-1.04+ (0.57)	0.04 (0.54)
Exploit II X Prior Success				0.10* (0.04)	-0.02 (0.03)
Explore X Prior Success				0.78 (0.55)	-0.35 (0.22)
Constant	-4.39*** (0.34)	-4.38*** (0.34)	-4.37*** (0.34)	-4.34*** (0.34)	-4.45*** (0.34)
<i>N</i>	12700	12699	12777	12700	12699
$\chi^2$	115	115	264	133	303
Prob > $\chi^2$	0.000	0.000	0.000	0.000	0.000
Log Likelihood	-1871	-1870	-1783	-1873	-1766

Standard errors in parentheses. +  $p < 0.10$ , \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ .  
Clustered by Assignee. Year and Technology Class Fixed Effects Included

**Table A11: Complementary Log-Log Model for Probability of Failure with Alternate Exploit Measure**

	(F1)	(F2)	(F3)	(F4)	(F5)
Explore	-0.06 (0.11)	-0.06 (0.11)			0.01 (0.11)
Exploit II	-0.02 (0.02)	-0.02 (0.02)			-0.03 (0.03)
Explore X Exploit II		-0.01 (0.02)			-0.01 (0.02)
Prior Success Experience			-0.24*** (0.05)		-0.26*** (0.05)
Prior Failure Experience			0.15** (0.05)		0.15* (0.06)
Exploit II X Prior Failure				0.02 (0.01)	0.01 (0.02)
Explore X Prior Failure				0.17 (0.16)	0.11 (0.16)
Exploit II X Prior Success				-0.01 (0.02)	0.04 (0.02)
Explore X Prior Success				-0.03 (0.15)	0.18 (0.17)
Constant	-2.93*** (0.15)	-2.93*** (0.15)	-3.00*** (0.15)	-2.94*** (0.15)	-3.01*** (0.15)
<i>N</i>	13386	13385	13464	13386	13385
$\chi^2$	508	513	523	506	538
Prob > $\chi^2$	0.000	0.000	0.000	0.000	0.000
Log Likelihood	-4015	-4013	-4037	-4015	-3994

Standard errors in parentheses. \*  $p < 0.10$ , \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ .  
Clustered by Assignee. Year and Technology Class Fixed Effects Included

**Table A12: Complementary Log-Log Model for Probability of Success and Failure with Square Root Discount Factor for Experience**

	Probability of Success					Probability of Failure				
	(S1)	(S2)	(S3)	(S4)	(S5)	(F1)	(F2)	(F3)	(F4)	(F5)
Explore	0.34*** (0.08)	0.35*** (0.08)			0.37*** (0.08)	-0.04 <sup>+</sup> (0.02)	-0.07** (0.02)			-0.07* (0.03)
Exploit	-0.49*** (0.05)	-0.4*** (0.05)			-0.49*** (0.05)	0.33*** (0.03)	0.35*** (0.03)			0.35*** (0.03)
Explore X Exploit		0.01 (0.04)			0.02 (0.04)		0.11*** (0.02)			0.11*** (0.02)
Prior Success Experience			0.36*** (0.05)		0.38*** (0.05)			-0.14*** (0.04)		-0.14*** (0.03)
Prior Failure Experience			-0.42*** (0.08)		-0.41*** (0.09)			0.05 (0.03)		0.07* (0.03)
Exploit X Prior Failure				0.21** (0.06)	0.02 (0.05)				0.03 (0.03)	0.01 (0.03)
Explore X Prior Failure				-0.10 (0.07)	-0.01 (0.07)				0.01 (0.02)	0.01 (0.02)
Exploit X Prior Success				-0.16* (0.07)	0.06 (0.04)				-0.04 (0.04)	-0.03 (0.04)
Explore X Prior Success				0.01 (0.10)	-0.101* (0.04)				0.01 (0.03)	0.02 (0.03)
Constant	-4.21*** (0.34)	-4.21*** (0.34)	-4.37*** (0.34)	-4.38*** (0.35)	-4.25*** (0.34)	-3.15*** (0.15)	-3.16*** (0.15)	-2.98*** (0.15)	-2.94*** (0.15)	-3.20*** (0.15)
<i>N</i>	12699	12699	12777	12699	12699	13385	13385	13464	13385	13385
$\chi^2$	186	185	126	85	329	606	620	526	509	657
Prob > $\chi^2$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Log Likelihood	-1823	-1823	-1852	-1874	-1785	-3950	-3943	-4047	-4012	-3935

**Table A13: Complementary Log-Log Model for Probability of Success and Failure with Linear Discount Factor for Experience**

	(S1)	(S2)	(S3)	(S4)	(S5)	(F1)	(F2)	(F3)	(F4)	(F5)
Explore	0.34*** (0.08)	0.35*** (0.08)			0.37*** (0.08)	-0.04+ (0.02)	-0.07** (0.02)			-0.07* (0.03)
Exploit	-0.49*** (0.05)	-0.49*** (0.05)			-0.49*** (0.05)	0.33*** (0.03)	0.35*** (0.03)			0.35*** (0.03)
Explore X Exploit		0.01 (0.04)			0.02 (0.04)		0.11*** (0.02)			0.11*** (0.02)
Prior Success Experience			0.35*** (0.05)		0.37*** (0.05)			-0.14*** (0.04)		-0.13*** (0.03)
Prior Failure Experience			-0.40*** (0.07)		-0.39*** (0.08)			0.05 (0.03)		0.06* (0.02)
Exploit X Prior Failure				0.19** (0.06)	0.01 (0.05)				0.03 (0.03)	0.01 (0.03)
Explore X Prior Failure				-0.10 (0.07)	-0.01 (0.07)				0.01 (0.02)	0.02 (0.02)
Exploit X Prior Success				-0.15* (0.07)	0.06+ (0.03)				-0.04 (0.04)	-0.03 (0.04)
Explore X Prior Success				0.01 (0.10)	-0.09* (0.04)				-0.01 (0.03)	0.01 (0.03)
Constant	-4.21*** (0.34)	-4.21*** (0.34)	-4.39*** (0.34)	-4.37*** (0.35)	-4.27*** (0.3)	-3.15*** (0.15)	-3.16*** (0.15)	-2.98*** (0.15)	-2.94*** (0.15)	-3.20*** (0.15)
$N$	12699	12699	12777	12699	12699	13385	13385	13464	13385	13385
$\chi^2$	186	185	140	83	349	606	620	526	510	658
Prob > $\chi^2$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Log Likelihood	-1823	-1823	-1852	-1875	-1783	-3950	-3943	-4047	-4012	-3935



## APPENDIX B

**Table B1: Model Notation**

$t$	Time, $t \in [0, T]$ : 0 (T) denotes the start (end) of the planning horizon.
$i(t)$	Rate of exploration efforts at time $t$ , $i(t) \geq 0$ ; control variable.
$e(t)$	Rate of exploitation efforts at time $t$ , $e(t) \geq 0$ ; control variable.
$\mu(t)$	Mean technical performance at time $t$ , $\mu(t) \geq 0$ ; $\mu(0)$ given; state variable.
$\sigma(t)$	Variance technical performance at time $t$ , $\sigma(t) \geq 0$ ; $\sigma(0)$ given; state variable.
$\alpha_0 (\alpha_1)$	Marginal impact of exploration on the mean (variance); $\alpha_0 (\alpha_1) > 0$ .
$\beta_0 (\beta_1)$	Marginal impact of exploitation on the mean (variance); $\beta_0 (\beta_1) > 0$ .
$c_0$	Parameter indicating marginal cost of exploration.
$c_1$	Parameter indicating marginal cost of exploitation.
$j=S, L$	Managers short-term and long-term objectives, respectively
$Z(t)$	Random variable indicating the innovation's technical performance at time $t$ ; $Z(t) \sim N(\mu(t), \sigma(t))$ .
$w_j$	Marginal value of a unit increase in technical performance
$P_j$	Normalized statistic such that $\Phi^{-1}(P_j) = \text{Probability}\{Z(t) \leq z_p(t)\}$ .
$P_j > 0 (< 0)$	Denotes risk seeking (risk averse) manager
$\lambda_0(t)$	Instantaneous marginal value of a unit increase in mean technical performance at time $t$ .
$\lambda_1(t)$	Instantaneous marginal value of a unit increase in the variance technical performance at time $t$ .
$x_0(t)$	Cumulative distributed marginal value of a unit increase in mean technical performance at time $t$ .
$x_1(t)$	Cumulative distributed marginal value of a unit increase in the variance technical performance at time $t$ .

## Hamiltonian

The Hamiltonian function,  $H$ , to be maximized is provided below. The first terms in Equations (1) and (2) are transformed into the reverse time expression, and represent the marginal value of a unit change in the mean and variance at time  $t$ , which results from a unit of exploration invested at time  $\tau$ .

$$\begin{aligned}
 H = & w_S z_S - 1/2 c_0 i(t)^2 - 1/2 c_1 e(t)^2 \\
 & + \alpha_0 i(t) \mu(t) \int_t^T \theta(\tau-t) \lambda_0(\tau) d\tau + \lambda_0(t) \beta_0 e(t) \mu(t) \\
 & + \alpha_1 i(t) \sigma(t) \int_t^T \theta(\tau-t) \lambda_1(\tau) d\tau - \lambda_1(t) \beta_1 e(t) \sigma(t)
 \end{aligned} \tag{B1}$$

## Optimality Conditions

Optimality Conditions for the rates of exploration and exploitation are as follows:

$$\partial H / \partial e(t) = 0, \quad e(t) \geq 0 \tag{B2}$$

$$\partial H / \partial i(t) = 0, \quad i(t) \geq 0 \tag{B3}$$

The necessary and sufficient conditions for optimality are as follows:

$$d\lambda_0(t)/dt = -\partial H / \partial \mu(t) \quad \text{and} \quad \lambda_0(T) = w_L \tag{B4}$$

$$d\lambda_1(t)/dt = -\partial H / \partial \sigma(t) \quad \text{and} \quad \lambda_1(T) = w_L P_L \tag{B5}$$

**PROOF OF THEOREM 1.** The results of Theorem 1 follow from Equations (B4) and (B5).

## PROOF OF COROLLARY 1a.

To prove that  $\lambda_0(t) > 0$  and  $d\lambda_0(t)/dt < 0$ . From Theorem 1(i) we obtain:

$$\frac{d\lambda_0(t)}{dt} = -w_S - \alpha_0 i(t) \int_t^T \theta(\tau-t) \lambda_0(\tau) d\tau - \lambda_0(t) \beta_0 e(t) \tag{B6}$$

(Proof by Contradiction) To prove that  $\lambda_0(t) > 0 \quad \forall t$ , suppose there exists some  $t_0 \in [0, T]$  such that  $\lambda_0(t_0) < 0$ . Since  $\lambda_0(T) > 0$  it follows that there exists  $\lambda'_0(t_1) > 0$  for some  $t_1 \in (t_0, T)$ . From Theorem 1(i) since  $w_S, \alpha_0, i(t_1), \beta_0, e(t_1)$ , and  $\theta(\tau-t) > 0$  for  $\tau \geq t$  then  $\lambda_0(t_1) < 0$  must

hold. Again, since  $\lambda_0(T) > 0$  it follows that there exists  $d\lambda_0(t_2)/dt > 0$  for some  $t_2 \in (t_1, T)$  from which we obtain  $\lambda_0(t_2) < 0$ . Repeating the process for a sufficiently large  $n$ ,  $t_n$  converges to  $T$  such that  $\lambda_0(t_n) = \lambda_0(T) < 0$ . However, since we know  $\lambda_0(T) > 0$ , it has been shown by contradiction that  $\lambda_0(t) < 0$  is not possible so that  $\lambda_0(t) > 0 \forall t$ . By similar reasoning, from Equation (5) given  $x_0(t) = \int_t^T \theta(\tau-t) \lambda_0(\tau) d\tau$  and  $x_0(T) = 0$  it follows that  $x_0(t) > 0 \forall t$  must hold. Since  $\lambda_0(t) > 0$ ,  $w_s$ ,  $\alpha_0$ ,  $i(t)$ ,  $\beta_0$ ,  $e(t) > 0$  and  $\theta(\tau-t) > 0$  for  $\tau \geq t$  we obtain  $d\lambda_0(t)/dt < 0 \forall t$ .

To complete the proof we need to show that  $dx_0(t)/dt < 0$  given  $\lambda_0(t) > 0$  and  $d\lambda_0(t)/dt < 0$ .

$$\text{Given } x_0(t) = \int_t^T \theta(\tau-t) \lambda_0(\tau) d\tau \text{ we obtain } \frac{dx_0(t)}{dt} = \int_t^T \frac{d\theta(\tau-t)}{dt} \lambda_0(\tau) d\tau - \theta(0) \lambda_0(t). \quad (\text{B7})$$

Let  $y = \tau - t$  therefore  $\theta(\tau-t) = \theta(y)$  and it follows that  $\frac{d\theta(\tau-t)}{dt} = \frac{-d\theta(\tau-t)}{\tau}$ .

$$\text{From Equation (B7) we obtain } \frac{dx_0(t)}{dt} = - \int_t^T \frac{d\theta(\tau-t)}{d\tau} \lambda_0(\tau) d\tau - \theta(0) \lambda_0(t). \quad (\text{B8})$$

Integrating by parts the first term of Equation (B8) we obtain:

$$\int_t^T \frac{d\theta(\tau-t)}{d\tau} \lambda_0(\tau) d\tau = \theta(\tau-t) \lambda_0(\tau) \Big|_t^T - \int_t^T \frac{d\lambda_0(\tau-t)}{d\tau} \theta(\tau-t) d\tau$$

Expanding terms above we obtain:

$$\int_t^T \frac{d\theta(\tau-t)}{d\tau} \lambda_0(\tau) d\tau = \theta(T-t) \lambda_0(T) - \theta(0) \lambda_0(t) - \int_t^T \frac{d\lambda_0(\tau-t)}{d\tau} \theta(y) dy \quad (\text{B9})$$

Substituting (B9) into (B7) we obtain:

$$\begin{aligned} \frac{dx_0(t)}{dt} &= - \left[ \theta(T-t) \lambda_0(T) - \theta(0) \lambda_0(t) - \int_t^T \frac{d\lambda_0(y)}{dy} \theta(y) dy \right] - \theta(0) \lambda_0(t) \\ \frac{dx_0(t)}{dt} &= -\theta(T-t) \lambda_0(T) + \int_t^T \frac{d\lambda_0(\tau)}{d\tau} \theta(y) dy \end{aligned} \quad (\text{B10})$$

From Equation (B10) since  $\theta(T-t)$ ,  $\theta(y)$ ,  $\lambda_0(T) \geq 0$  and  $\frac{d\lambda_0(\tau)}{d\tau} \leq 0$ , we know the sum on the right hand side of Equation (B10) is negative, so that  $\frac{dx_0(t)}{dt} \leq 0$ . QED

### PROOF OF COROLLARY 1b.

We consider all the possible optimal solutions for  $\lambda_1(t)$  and  $d\lambda_1(t)/dt$  for the four possible terminal conditions as follows:

- Case I which assumes that  $\lambda_1(T)<0$  and  $d\lambda_1(T)/dt>0$  hold,
- Case II which assumes that  $\lambda_1(T)>0$  and  $d\lambda_1(T)/dt<0$  hold,
- Case III which assumes that  $\lambda_1(T)<0$  and  $d\lambda_1(T)/dt<0$  hold.
- Case IV which assumes that  $\lambda_1(T)>0$  and  $d\lambda_1(T)/dt>0$  hold.

#### Case I. Assume $\lambda_1^*(T)<0$ and $d\lambda_1^*(T)/dt>0$ hold.


From Theorem 1(i) we have  $\lambda_1(T)=w_L P_L$ . Therefore, given  $w_L>0$ , for  $\lambda_1(T)<0$  then  $P_L<0$  must hold. From Equation (6) we obtain  $x_1(T)=0$ , so at the terminal time, Theorem 1(i)

simplifies to the following: 
$$\frac{d\lambda_1(T)}{dt} = -w_S P_S + \lambda_1(T) \beta_1 e(T) \quad (B11)$$

From Equation (B11), to obtain  $d\lambda_1(T)/dt>0$ , given  $\lambda_1(T)<0$ , we must have  $P_S<0$ . In conclusion, we have shown that  $P_L<0$  and  $P_S<0$  must hold under the suppositions in Case I,  $\lambda_1(T)<0$  and  $d\lambda_1(T)/dt>0$ . This gives us three possible solutions for  $\lambda_1(t)$  and  $d\lambda_1(t)/dt$ , as described below in 1-3 and illustrated in Figure B1.

1.  $\lambda_1(t)<0$ ;  $d\lambda_1(t)/dt>0 \forall t$ ,
2.  $\lambda_1(t)<0$ ,  $d\lambda_1(t)/dt\leq 0$ ,  $t\in[0,t_1]$ ;  $\lambda_1(t)<0$ ,  $d\lambda_1(t)/dt>0$ ,  $t\in(t_1,T]$  or
3.  $\lambda_1(t)\geq 0$ ,  $d\lambda_1(t)/dt<0$   $t\in[0,t_1]$ ;  $\lambda_1(t)<0$ ,  $d\lambda_1(t)/dt\leq 0$ ,  $t\in(t_1,t_2]$ ;  $\lambda_1(t)<0$ ,  $d\lambda_1(t)/dt>0$ ,  $t\in(t_2,T]$ .

Below we consider the feasibility of these three possible solutions.



**Figure B1: Case I  $\lambda_1(T) < 0, d\lambda_1(T)/dt > 0$**

**Case I.1:**  $\lambda_1(t) < 0; d\lambda_1(t)/dt > 0 \forall t$

From Theorem 1(i) we know:

$$\frac{d\lambda_1(t)}{dt} = -w_s P_s - \alpha_1 i(t) \int_t^T \theta(\tau-t) \lambda_1(\tau) d\tau + \lambda_1(t) \beta_1 e(t) \quad (B12)$$

Applying the reverse time transformation to the sixth term of Equation (B1) and rearranging for  $i(t)$  gives us:

$$i(t) = \frac{\lambda_1(t) \int_0^t \alpha_1 \theta(t-\tau) i(\tau) \sigma(\tau) d\tau}{\alpha_1 \sigma(t) \int_t^T \theta(\tau-t) \lambda_1(\tau) d\tau} \quad (B13)$$

Substituting Equations (2) and (B13) into Equation (B12) we obtain:

$$\frac{d\lambda_1(t)}{dt} = -w_s P_s - \lambda_1(t) \left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] \quad (B14)$$

From Equation (B14) applying backward integration we obtain:

$$\lambda_1(t) = w_L P_L + \int_t^T \left[ w_s P_s + \lambda_1(y) \left[ \frac{d\sigma(y)}{dy} / \sigma(y) \right] \right] dy \quad (B15)$$

Based on Equation (B14), since we know  $P_s < 0$  must hold, we consider two possible scenarios:

(a)  $P_s < 0$  and  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] > 0$  for  $t \in [0, T)$ ,

(b)  $P_S < 0$  and  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] < 0$  for  $t \in [0, T]$ .

For scenario (a), given  $\lambda_1(t) < 0$  and the above conditions, we know  $d\lambda_1(t)/dt > 0$  must hold. Therefore, Case III.1 is feasible under scenario (a). However, since  $\sigma(t)$  is increasing in  $i(t)$  and decreasing in  $e(t)$ , and since  $i(t)$  is increasing in  $\lambda_1(t)$ , while  $e(t)$  is decreasing in  $\lambda_1(t)$ , then if  $\lambda_1(t) < 0$  holds, it is not likely that  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] > 0$  also holds.

Therefore, Case I.1, while feasible under scenario (a), is unlikely to occur. In addition, given  $\lambda_1(t) < 0$  and the above conditions, Case I.1 is feasible under scenario (b) when  $|P_S|$  is large.

**Case I. 2:**  $\lambda_1(t) < 0, d\lambda_1(t)/dt \leq 0, t \in [0, t_1]; \lambda_1(t) < 0, d\lambda_1(t)/dt > 0, t \in (t_1, T]$

Based on the analysis of Case I.1, we know the conditions under which  $t \in (t_1, T]$  in Case I.2 hold. Since we know  $P_S < 0$  must hold, based on (B14) we consider two possible scenarios:

(a)  $P_S < 0$  and  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] > 0$  for  $t \in [0, t_1]$ ,

(b)  $P_S < 0$  and  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] < 0$  for  $t \in [0, t_1]$ .

For scenario (a), given  $\lambda_1(t) < 0$  and the above conditions, we know  $d\lambda_1(t)/dt < 0$  cannot occur. Therefore, Case I.2 is infeasible under scenario (a). In addition, given  $\lambda_1(t) < 0$  and the above conditions, Case I.2 is feasible under scenario (b) when  $|P_S|$  is small. However, based on the analysis of Case I.1, we know the conditions under which  $t \in [t_1, T]$  hold are more likely when  $|P_S|$  is large. Therefore, Case I.2, while feasible under scenario (b), is unlikely to occur.

**Case I. 3:**  $\lambda_1(t) \geq 0, d\lambda_1(t)/dt < 0, t \in [0, t_1]; \lambda_1(t) < 0, d\lambda_1(t)/dt \leq 0, t \in (t_1, t_2]; \lambda_1(t) < 0, d\lambda_1(t)/dt > 0, t \in (t_2, T]$ .

First, corresponding to Cases I.1 and I.2, we know the conditions under which  $t \in (t_1, T]$  from Case I.3 hold. Second, at  $t = t_1$  we know  $\lambda_1(t_1) = 0$  and  $d\lambda_1(t_1)/dt < 0$  hold. From Equation (B14) since we know  $P_S < 0$  and  $P_L < 0$  and then  $\lambda_1(t_1) = 0$  and  $d\lambda_1(t_1)/dt < 0$  cannot occur and Case I.3 is infeasible.

Focusing on Case I.1, we now consider the solutions for  $x_1(t)$  and  $dx_1(t)/dt$  when  $\lambda_1(t) < 0$ ;  $d\lambda_1(t)/dt > 0 \forall t$ . The proof that  $x_1(t) < 0$  and  $dx_1(t)/dt > 0$ , given  $\lambda_1(t) < 0$  and  $d\lambda_1(t)/dt > 0$ , is analogous to the proof that  $x_0(t) > 0$  and  $dx_0(t)/dt < 0$  must hold given  $\lambda_0(t) > 0$  and  $d\lambda_0(t)/dt < 0$  (see Corollary 1a).

Lastly, we need to prove that  $x_1(t) > \lambda_1(t)$ . Since  $\lambda_1(t) < 0$  and  $d\lambda_1(t)/dt \geq 0$ , we know

$$0 > \lambda_1(\tau) \geq \lambda_1(t) \text{ for } \tau \geq t. \text{ This gives us } \int_t^T \theta(\tau-t) \lambda_1(\tau) d\tau > \int_t^T \theta(\tau-t) \lambda_1(t) d\tau. \quad (\text{B16})$$

$$\text{Since } \int_t^\infty \theta(\tau-t) d\tau = 1 \text{ and } \lambda_1(t) < 0, \text{ we know } \int_t^T \theta(\tau-t) \lambda_1(t) d\tau > \int_t^\infty \theta(\tau-t) \lambda_1(t) d\tau = \lambda_1(t). \quad (\text{B17})$$

Comparing  $x_1(t)$  from Equation (6) with Equations (B16) and (B17) above, we obtain  $x_1(t) > \lambda_1(t)$ , as desired.


In summary we have shown that Corollary 1b.I is most likely to occur given the conditions assumed.

**Case II. Assume  $\lambda_1(T) > 0$  and  $d\lambda_1(T)/dt < 0$  hold.**

From Theorem 1(i) we have  $\lambda_1(T) = w_L P_L$ . Therefore, given  $w_L > 0$ , for  $\lambda_1(T) > 0$  to hold then  $P_L > 0$  must hold. From Equation (B11), to obtain  $d\lambda_1(T)/dt < 0$ , given  $\lambda_1(T) > 0$ , we must have  $P_S > 0$ . In conclusion, we have shown that  $P_L > 0$  and  $P_S > 0$  hold under the suppositions in Case II,  $\lambda_1(T) > 0$  and  $d\lambda_1(T)/dt < 0$ . This gives us three possible solutions for  $\lambda_1(t)$  and  $d\lambda_1(t)/dt$ , as described below in 1-3 and illustrated in Figure A2.

1.  $\lambda_1(t) > 0$ ;  $d\lambda_1(t)/dt < 0 \forall t \in [0, T]$ ,
2.  $\lambda_1(t) > 0$ ,  $d\lambda_1(t)/dt \geq 0$ ,  $t \in [0, t_1]$ ;  $\lambda_1(t) > 0$ ,  $d\lambda_1(t)/dt < 0$ ,  $t \in (t_1, T]$  or
3.  $\lambda_1(t) \leq 0$ ,  $d\lambda_1(t)/dt > 0 \forall t \in [0, t_1]$ ;  $\lambda_1(t) > 0$ ,  $d\lambda_1(t)/dt \geq 0$ ,  $t \in (t_1, t_2]$ ;  $\lambda_1(t) > 0$ ,  $d\lambda_1(t)/dt < 0$ ,  $t \in (t_2, T]$ .

Below we consider the feasibility of these three possible solutions.



**Figure B2: Case II  $\lambda_1(T) > 0$ ,  $d\lambda_1(T)/dt < 0$**

**Case II.1:**  $\lambda_1(t) > 0$ ;  $d\lambda_1(t)/dt < 0$  for  $t \in [0, T]$ ,

Based on Equation (B14), we consider two possible scenarios since we know  $P_S > 0$  must hold:

(a)  $P_S > 0$  and  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] > 0$  for  $t \in [0, T]$ ,

(b)  $P_S > 0$  and  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] < 0$  for  $t \in [0, T]$ .

For scenario (a), given  $\lambda_1(t) > 0$  and the above conditions, we know  $d\lambda_1(t)/dt < 0$  must hold. Therefore, Case II.1 is feasible under scenario (a). Similarly, given  $\lambda_1(t) > 0$  and the above conditions, Case II.1 is feasible under scenario (b) when  $|P_S|$  is large.

**Case II.2:**  $\lambda_1(t) > 0$ ,  $d\lambda_1(t)/dt \geq 0$ ,  $t \in [0, t_1]$ ;  $\lambda_1(t) > 0$ ,  $d\lambda_1(t)/dt < 0$ ,  $t \in (t_1, T]$

Based on the analysis of Case II.1, we know the conditions under which  $t \in (t_1, T]$  in Case II.2 hold. With  $P_S > 0$ , we consider two possible scenarios:

(a)  $P_S > 0$  and  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] > 0$  for  $t \in [0, t_1]$ ,

(b)  $P_S > 0$  and  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] < 0$  for  $t \in [0, t_1]$ .

For scenario (a), given  $\lambda_1(t) > 0$  and the above conditions, we know  $d\lambda_1(t)/dt > 0$  cannot occur. Therefore, Case II.2 is infeasible under scenario (a). In addition, given  $\lambda_1(t) > 0$  and the above conditions, Case II.2 is feasible under scenario (b) when  $|P_S|$  is small.



However, note that  $\sigma(t)$  is increasing in  $i(t)$  and decreasing in  $e(t)$ ;  $i(t)$  is increasing in  $\lambda_1(t)$ ; while  $e(t)$  is decreasing in  $\lambda_1(t)$ . Given  $\lambda_1(t) > 0$  holds, it is not likely that  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] < 0$  also holds. Therefore, Case II.2, while feasible, is unlikely to occur.

**Case II.3:**  $\lambda_1(t) \leq 0$ ,  $d\lambda_1(t)/dt > 0$   $t \in [0, t_1]$ ;  $\lambda_1(t) > 0$ ,  $d\lambda_1(t)/dt \geq 0$ ,  $t \in (t_1, t_2]$ ;  $\lambda_1(t) > 0$ ,  $d\lambda_1(t)/dt < 0$ ,  $t \in (t_2, T]$ .

First, corresponding to Cases II.1 and II.2, we know the conditions under which  $t \in (t_1, T]$  from Case II.3 hold. Second, at  $t = t_1$  we know  $\lambda_1(t_1) = 0$  and  $d\lambda_1(t_1)/dt > 0$ . With  $P_S > 0$  and  $P_L > 0$ , then from Equation (B15) Case II.3 is infeasible.

Focusing on Case II.1, which is most likely, we now consider the solutions for  $x_1(t)$  and  $dx_1(t)/dt$ . The proof that  $x_1(t) > 0$  and  $dx_1(t)/dt < 0$ , given  $\lambda_1(t) > 0$  and  $d\lambda_1(t)/dt < 0$ , is analogous to the proof that  $x_0(t) > 0$  and  $dx_0(t)/dt < 0$  must hold given  $\lambda_0(t) > 0$  and  $d\lambda_0(t)/dt < 0$  (see Corollary 1a). The proof that  $x_1(t) \leq \lambda_1(t)$  when  $\lambda_1(t) > 0$  and  $d\lambda_1(t)/dt \leq 0$  is analogous to the proof that  $x_1(t) \geq \lambda_1(t)$  when  $\lambda_1(t) < 0$  and  $d\lambda_1(t)/dt \geq 0$  hold (see Corollary 1b.I)

In summary we have shown that Corollary 1b.II is most likely to occur given the conditions assumed for Case II.

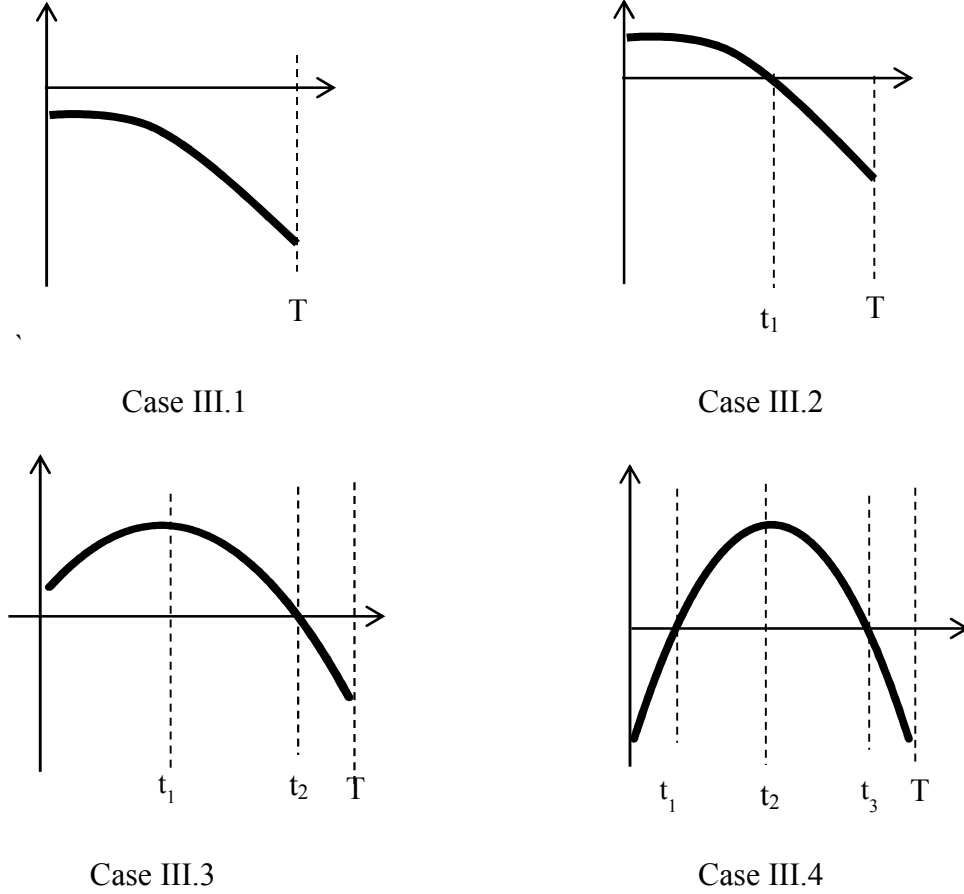
**Case III. Assume  $\lambda_1^*(T) < 0$  and  $d\lambda_1^*(T)/dt < 0$  hold.**

From Theorem 1(i) we have  $\lambda_1(T) = w_L P_L$ . Therefore, given  $w_L > 0$ , for  $\lambda_1(T) < 0$  to hold then  $P_L < 0$  must hold. From Equation (B11) for  $d\lambda_1(T)/dt < 0$  and  $\lambda_1(T) < 0$ , then either  $P_S > 0$  or  $P_S < 0$  hold. This gives us four possible solutions for  $\lambda_1(t)$  and  $d\lambda_1(t)/dt$ , as described below in 1-4 and illustrated in Figure B2.

1.  $\lambda_1(t) < 0$ ;  $d\lambda_1(t)/dt < 0 \forall t \in [0, T]$ .
2.  $\lambda_1(t) \geq 0$ ,  $d\lambda_1(t)/dt < 0$  for  $t \in [0, t_1]$ ;  $\lambda_1(t) < 0$ ,  $d\lambda_1(t)/dt < 0$ ,  $t \in (t_1, T]$ .
3.  $\lambda_1(t) > 0$ ,  $d\lambda_1(t)/dt \geq 0$   $t \in [0, t_1]$ ;  $\lambda_1(t) \geq 0$ ,  $d\lambda_1(t)/dt < 0$ ,  $t \in (t_1, t_2]$ ;  $\lambda_1(t) < 0$ ,  $d\lambda_1(t)/dt < 0$ ,  $t \in (t_2, T]$ .

4.  $\lambda_1(t) \leq 0, d\lambda_1(t)/dt > 0 \ t \in [0, t_1]$ ;  $\lambda_1(t) > 0, d\lambda_1(t)/dt \geq 0, t \in (t_1, t_2]$ ;  $\lambda_1(t) > 0, d\lambda_1(t)/dt < 0, t \in (t_2, t_3]$ ;  $\lambda_1(t) < 0, d\lambda_1(t)/dt > 0$  for  $t \in (t_3, T]$ .

Below we consider the feasibility of these four possible solutions.



**Figure B3: Case III  $\lambda_1(T) < 0, d\lambda_1(T)/dt < 0$**

**Case III.1:**  $\lambda_1(t) < 0; d\lambda_1(t)/dt < 0 \ \forall t \in [0, T]$ ,

Based on Equation (B14), we consider four possible scenarios based on the two possible solutions for  $P_S > 0$  and  $P_S < 0$ :

- (a)  $P_S > 0$  and  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] > 0$  for all  $t \in [0, T]$ ,
- (b)  $P_S > 0$  and  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] < 0$  for all  $t \in [0, T]$ ,

(c)  $P_S < 0$  and  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] > 0$  for all  $t \in [0, T]$ ,

(d)  $P_S < 0$  and  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] < 0$  for all  $t \in [0, T]$ .

For scenario (a), given  $\lambda_1(t) < 0$  and the above conditions, Case III.1 is feasible when  $|P_S|$  is large. However, since  $\sigma(t)$  is increasing in  $i(t)$  and decreasing in  $e(t)$ , and since  $i(t)$  is increasing in  $\lambda_1(t)$ , while  $e(t)$  is decreasing in  $\lambda_1(t)$ , then if  $\lambda_1(t) < 0$  holds, it is not likely that  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] > 0$  also holds. Therefore, Case III.1, while feasible under scenario (a), is unlikely to occur. For scenario (b), given  $\lambda_1(t) < 0$  and the above conditions, we know  $d\lambda_1(t)/dt < 0$  must hold. Therefore, Case III.1 is feasible under scenario (b). However, for scenario (c), given  $\lambda_1(t) < 0$  and the above conditions, we know  $d\lambda_1(t)/dt < 0$  cannot occur. Therefore, Case III.1 is infeasible under scenario (c). Lastly, for scenario (d), given  $\lambda_1(t) < 0$  and the above conditions, Case III.1 is feasible when  $|P_S|$  is small.

**Case III.2:**  $\lambda_1(t) \geq 0$ ,  $d\lambda_1(t)/dt < 0$  for  $t \in [0, t_1]$ ;  $\lambda_1(t) < 0$ ,  $d\lambda_1(t)/dt < 0$ ,  $t \in (t_1, T]$

Based on the analysis of Case III.1, we know the conditions under which  $t \in (t_1, T]$  in Case III.2 hold. Also, at  $t = t_1$  we know  $\lambda_1(t_1) = 0$  and  $d\lambda_1(t_1)/dt < 0$ . Since we know  $P_L < 0$ , then from Equation (B14), we know  $\lambda_1(t_1) = 0$ ,  $d\lambda_1(t_1)/dt > 0$  is infeasible when  $P_S < 0$ . Therefore,  $P_S > 0$  must hold. To prove the conditions for feasibility of  $\lambda_1(t) > 0$  and  $d\lambda_1(t)/dt < 0$  for  $t \in [0, t_1]$ , given  $\lambda_1(T) = w_L P_L < 0$  and  $P_S > 0$  must hold, we consider two possible scenarios:

(a)  $P_S > 0$  and  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] > 0$  for  $t \in [t_1, T]$ ,

(b)  $P_S > 0$  and  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] < 0$  for  $t \in [t_1, T]$ .

For scenario (a), given  $\lambda_1(t) > 0$  and the above conditions, we know  $d\lambda_1(t)/dt < 0$  must hold. Therefore, Case III.2 is feasible under scenario (a). Also, given  $\lambda_1(t) > 0$  and the above conditions, Case III.1 is feasible under scenario (b) when  $|P_S|$  is large. However, since  $\sigma(t)$  is increasing in  $i(t)$  and decreasing in  $e(t)$ , and since  $i(t)$  is increasing in  $\lambda_1(t)$ ,

while  $e(t)$  is decreasing in  $\lambda_1(t)$ , then if  $\lambda_1(t) > 0$  holds, it is not likely that  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] < 0$  also holds. Therefore, Case III.2, while feasible under scenario (b), is unlikely to occur.

**Case III.3:**  $\lambda_1(t) > 0, d\lambda_1(t)/dt \geq 0, t \in [0, t_1]$ ;  $\lambda_1(t) \geq 0, d\lambda_1(t)/dt < 0, t \in (t_1, t_2]$ ;  $\lambda_1(t) < 0, d\lambda_1(t)/dt < 0, t \in (t_2, T]$ .

First, corresponding to Cases III.1 and III.2, we know the conditions under which  $t \in (t_1, T]$  from Case III.3 hold. Therefore  $P_S > 0$  must hold. This gives us two possible scenarios:

(a)  $P_S > 0$  and  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] > 0$  for  $t \in [0, t_1]$ ,

(b)  $P_S > 0$  and  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] < 0$  for  $t \in [0, t_1]$ .

For scenario (a), since at  $t = t_1$  we know  $\lambda_1(t_1) > 0$  and  $d\lambda_1(t_1)/dt = 0$ . From Equation (B14), for  $\lambda_1(t) > 0$  we know  $d\lambda_1(t_1)/dt = 0$  holds if and only if  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] < 0$ . Therefore, Case III.3 is infeasible under scenario (a). For scenario (b), given  $\lambda_1(t) > 0$  and the above conditions, Case III.3 is feasible when  $|P_S|$  is small. However, since  $\sigma(t)$  is increasing in  $i(t)$  and decreasing in  $e(t)$ , and since  $i(t)$  is increasing in  $\lambda_1(t)$ , while  $e(t)$  is decreasing in  $\lambda_1(t)$ , then if  $\lambda_1(t) > 0$  holds, it is not likely that  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] < 0$  also holds. Therefore, Case III.3, while feasible under scenario (b), is unlikely to occur.

**Case III.4:**  $\lambda_1(t) \leq 0, d\lambda_1(t)/dt > 0, t \in [0, t_1]$ ;  $\lambda_1(t) > 0, d\lambda_1(t)/dt \geq 0, t \in (t_1, t_2]$ ;  $\lambda_1(t) > 0, d\lambda_1(t)/dt < 0, t \in (t_2, t_3]$ ;  $\lambda_1(t) < 0, d\lambda_1(t)/dt > 0$  for  $t \in (t_3, T]$ .

Based on the analysis of Cases III.1, III.2 and III.3, we know the conditions under which  $t \in (t_1, T]$  in Case III.4 hold. Therefore,  $P_S > 0$  must hold. This gives us two possible scenarios:

(a)  $P_S > 0$  and  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] > 0$  for  $t \in [0, t_1]$ ,

(b)  $P_S > 0$  and  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] < 0$  for  $t \in [0, t_1]$ .

For scenario (a), given  $\lambda_1(t) < 0$  and the above conditions, Case III.4 is feasible when  $|P_S|$  is small. However, since  $\sigma(t)$  is increasing in  $i(t)$  and decreasing in  $e(t)$ , and since  $i(t)$  is increasing in  $\lambda_1(t)$ , while  $e(t)$  is decreasing in  $\lambda_1(t)$ , then if  $\lambda_1(t) < 0$  holds, it is not likely that  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] > 0$  also holds. Therefore, Case III.4, while feasible under scenario (b), is unlikely to occur. In addition, for scenario (b), given  $\lambda_1(t) < 0$  and the above conditions, we know  $d\lambda_1(t)/dt > 0$  cannot occur. Therefore, Case III.4 is infeasible under scenario (b).


In summary we have shown that Corollary 1b.III is likely to occur given the conditions assumed.

**Case IV. Assume  $\lambda_1^*(T) > 0$  and  $d\lambda_1^*(T)/dt > 0$  hold.**

From Theorem 1(i) we have  $\lambda_1(T) = w_L P_L$ . Therefore, given  $w_L > 0$ , for  $\lambda_1(T) > 0$  to hold then  $P_L > 0$  must hold. From Equation (B11) for  $d\lambda_1(T)/dt > 0$  to hold, given  $\lambda_1(T) > 0$ , then  $\lambda_1(T)\beta_1 e(T) > w_S P_S$  must hold. Since  $\lambda_1(T) = w_L P_L > 0$  then either  $P_S > 0$  or  $P_S < 0$  is feasible. This gives us five possible solutions for  $\lambda_1(t)$  and  $d\lambda_1(t)/dt$ , as described below in 1-5 and illustrated in Figure B4.

1.  $\lambda_1(t) > 0$ ;  $d\lambda_1(t)/dt > 0 \forall t \in [0, T]$ .
2.  $\lambda_1(t) > 0$ ,  $d\lambda_1(t)/dt \leq 0$ ,  $t \in [0, t_1]$ ;  $\lambda_1(t) > 0$ ,  $d\lambda_1(t)/dt > 0$ ,  $t \in (t_1, T]$
3.  $\lambda_1(t) \leq 0$ ,  $d\lambda_1(t)/dt > 0$ ,  $t \in [0, t_1]$ ;  $\lambda_1(t) > 0$ ,  $d\lambda_1(t)/dt > 0$ ,  $t \in (t_1, T]$
4.  $\lambda_1(t) < 0$ ,  $d\lambda_1(t)/dt \leq 0$ ,  $t \in [0, t_1]$ ;  $\lambda_1(t) \leq 0$ ,  $d\lambda_1(t)/dt > 0$ ,  $t \in (t_1, t_2]$ ;  $\lambda_1(t) > 0$ ,  $d\lambda_1(t)/dt > 0$ ,  $t \in (t_2, T]$ .
5.  $\lambda_1(t) \geq 0$ ,  $d\lambda_1(t)/dt < 0$   $t \in [0, t_1]$ ;  $\lambda_1(t) < 0$ ,  $d\lambda_1(t)/dt \leq 0$ ,  $t \in (t_1, t_2]$ ;  $\lambda_1(t) \leq 0$ ,  $d\lambda_1(t)/dt > 0$ ,  $t \in (t_2, t_3]$ ;  $\lambda_1(t) > 0$ ,  $d\lambda_1(t)/dt > 0$  for  $t \in (t_3, T]$ .

Below we consider the feasibility of these five possible solutions.



**Figure B4: Case IV  $\lambda_1(T) > 0$ ,  $d\lambda_1(T)/dt > 0$**

**Case IV.1:**  $\lambda_1(t) > 0$ ;  $d\lambda_1(t)/dt > 0 \forall t \in [0, T]$ ,

Based on Equation (B14), we consider four possible scenarios based on the two possible solutions for  $P_S > 0$  and  $P_S < 0$ :

- (a)  $P_S > 0$  and  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] > 0$  for all  $t \in [0, T]$ ,
- (b)  $P_S > 0$  and  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] < 0$  for all  $t \in [0, T]$ ,
- (c)  $P_S < 0$  and  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] > 0$  for all  $t \in [0, T]$ ,
- (d)  $P_S < 0$  and  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] < 0$  for all  $t \in [0, T]$ .

First, for scenario (a), given  $\lambda_1(t) > 0$  and the above conditions, we know  $d\lambda_1(t)/dt > 0$  cannot occur. Therefore, Case IV.1 is infeasible under scenario (a). Second, given  $\lambda_1(t) > 0$  and the above conditions, we know  $d\lambda_1(t)/dt > 0$  is feasible under scenario (b) when  $|P_S|$  is small. However, since  $\sigma(t)$  is increasing in  $i(t)$  and decreasing in  $e(t)$ , and since  $i(t)$  is increasing in  $\lambda_1(t)$ , while  $e(t)$  is decreasing in  $\lambda_1(t)$ , then if  $\lambda_1(t) > 0$  holds, it is not likely that  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] < 0$  also holds. Therefore, Case IV.1, while feasible under scenario (b), is unlikely to occur. Third, for scenario (c). Given  $\lambda_1(t) > 0$  and the above conditions, Case IV.1 is feasible under scenario (c) when  $|P_S|$  is large. Lastly, for scenario (d), given  $\lambda_1(t) > 0$  and the above conditions, we know  $d\lambda_1(t)/dt > 0$  must hold. Therefore, Case IV.1 is feasible under scenario (d).

**Case IV.2:**  $\lambda_1(t) > 0$ ,  $d\lambda_1(t)/dt \leq 0$  for  $t \in [0, t_1]$ ;  $\lambda_1(t) > 0$ ,  $d\lambda_1(t)/dt > 0$ ,  $t \in (t_1, T]$

Based on the analysis of Case IV.1, we know the conditions under which  $t \in (t_1, T]$  in Case IV.2 hold. Based on Equation (B14), we consider four possible scenarios based on the two possible solutions for  $P_S > 0$  and  $P_S < 0$ :

- (a)  $P_S > 0$  and  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] > 0$  for all  $t \in [0, t_1]$ ,
- (b)  $P_S > 0$  and  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] < 0$  for all  $t \in [0, t_1]$ ,
- (c)  $P_S < 0$  and  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] > 0$  for all  $t \in [0, t_1]$ ,
- (d)  $P_S < 0$  and  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] < 0$  for all  $t \in [0, t_1]$ .

For scenario (a), given  $\lambda_1(t) > 0$  and the above conditions, we know  $d\lambda_1(t)/dt < 0$  must hold. Therefore, Case IV.2 is feasible under scenario (a). Second, for scenario (b), given  $\lambda_1(t) > 0$  and the above conditions, Case IV.2 is feasible under scenario (b) when  $|P_S|$  is large. Third, for scenario (c), given  $\lambda_1(t) > 0$  and the above conditions, Case IV.2 is feasible under scenario (c) when  $|P_S|$  is small. Lastly, for scenario (d), given  $\lambda_1(t) > 0$  and the above conditions, we know  $d\lambda_1(t)/dt < 0$  cannot occur. Therefore, Case IV.2 is infeasible under scenario (d).

**Case IV.3:**  $\lambda_1(t) \leq 0, d\lambda_1(t)/dt > 0$  for  $t \in [0, t_1]$ ;  $\lambda_1(t) > 0, d\lambda_1(t)/dt > 0, t \in (t_1, T]$

Based on the analysis of Case IV.1, we know the conditions under which  $t \in (t_1, T]$  in Case IV.3 hold. Second, at  $t = t_1$  we know  $\lambda_1(t_1) = 0$  and  $d\lambda_1(t_1)/dt > 0$ . From Equation (B14), we know  $\lambda_1(t_1) = 0, d\lambda_1(t_1)/dt > 0$ , so that  $P_S < 0$  must hold. This gives us two possible scenarios:

(a)  $P_S < 0$  and  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] > 0$  for  $t \in [0, t_1]$ ,

(b)  $P_S < 0$  and  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] < 0$  for  $t \in [0, t_1]$ .

For scenario (a), given  $\lambda_1(t) < 0$  and the above conditions we know  $d\lambda_1(t)/dt > 0$  must hold. Therefore, Case IV.3 is feasible under scenario (a). In addition, for scenario (b), given  $\lambda_1(t) < 0$  and the above conditions, Case IV.3 is feasible under scenario (b), when  $|P_S|$  is large.

**Case IV.4:**  $\lambda_1(t) < 0, d\lambda_1(t)/dt \leq 0$   $t \in [0, t_1]$ ;  $\lambda_1(t) \leq 0, d\lambda_1(t)/dt > 0, t \in (t_1, t_2]$ ;  $\lambda_1(t) > 0, d\lambda_1(t)/dt > 0, t \in (t_2, T]$ .

Corresponding to Cases IV.1 and IV.3, we know the conditions under which the suppositions in Case IV.4 hold for  $t \in (t_1, T]$ . Therefore, we only consider  $P_S < 0$ . This gives us two possible scenarios:

(a)  $P_S < 0$  and  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] > 0$  for  $t \in [0, t_1]$ ,

(b)  $P_S < 0$  and  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] < 0$  for  $t \in [0, t_1]$ .

For scenario (a), given  $\lambda_1(t) < 0$  and the above conditions, we know  $d\lambda_1(t)/dt < 0$  cannot occur. Therefore, Case IV.4 is infeasible under scenario (a). For scenario (b), given  $\lambda_1(t) < 0$  and the above conditions, Case IV.4 is feasible when  $|P_S|$  is small. However, based on the analysis of Case IV.3, we know the conditions which hold during the interval  $t \in (t_1, T]$  are more likely when  $|P_S|$  is large, which contradicts Case IV.3. Therefore, Case IV.4, while feasible under scenario (b), is unlikely to occur.



**Case IV.5:**  $\lambda_1(t) \geq 0$ ,  $d\lambda_1(t)/dt < 0$   $t \in [0, t_1]$ ;  $\lambda_1(t) < 0$ ,  $d\lambda_1(t)/dt \leq 0$ ,  $t \in (t_1, t_2]$ ;  $\lambda_1(t) \leq 0$ ,  $d\lambda_1(t)/dt > 0$ ,  $t \in (t_2, t_3]$ ;  $\lambda_1(t) > 0$ ,  $d\lambda_1(t)/dt > 0$  for  $t \in (t_3, T]$ .

First, corresponding to Cases IV.1, IV.3 and IV.4, we know the conditions under which

Case IV.5 for  $t \in (t_1, T]$  hold. Therefore, we only consider  $P_S < 0$  and  $\left[ \frac{d\sigma(t)}{dt} / \sigma(t) \right] < 0$ .

Given  $\lambda_1(t) > 0$  we know  $d\lambda_1(t)/dt < 0$  cannot occur. Therefore, Case II.5 is infeasible.

In summary we have shown that Corollary 1b.IV is likely to occur given the conditions assumed.

**PROOF OF THEOREM 2.** The results of Theorem 2 follow from the optimality conditions in Equations (B2) and (B3). Note that if at some  $t \in [0, T]$  we obtain  $e(t) < 0$  for  $e(t)$  satisfying  $\partial H / \partial e(t) = 0$ , then  $e^*(t) = 0$  holds at that instant of time. Similarly, if we obtain  $i(t) < 0$  for  $i(t)$  satisfying  $\partial H / \partial i(t) = 0$ , then  $i^*(t) = 0$  holds at that instant of time.

**PROOF OF COROLLARY 2a.** The results of Corollary 2a are derived from Theorem 2. Consider  $e(t)$  such that it satisfies  $\partial H / \partial e(t) = 0$  (unconstrained). First, if  $e(t) > 0$  for all  $t \in [0, T]$ , then  $e^*(t) > 0$  over the entire planning horizon and no stopping or starting time occurs. Second, if  $t_0$  exists such that:  $e(t) < 0$  for  $t < t_0$ ,  $e(t_0) = 0$ , and  $e(t) > 0$  for  $t > t_0$ , then  $t_0$  is an optimal starting time for exploitation. Third, if  $t_0$  exists such that:  $e(t) > 0$  for  $t < t_0$ ,  $e(t_0) = 0$ , and  $e(t) < 0$  for  $t > t_0$ , then  $t_0$  is an optimal stopping time for exploitation. The analysis for starting and stopping times for  $i^*(t)$  is analogous and is omitted.

**PROOF OF COROLLARIES 2b and 2c.** The results of Corollaries 2b and 2c follow from Theorem 2.

**PROOF OF COROLLARY 3.** The results of Corollary 3 follow from Theorem 2.

**PROOF OF THEOREM 3.** The results of Theorem 3 follow from differentiation of the results from Theorem 2.

**PROOF OF COROLLARY 4a.** Recall from Equation (1) that

$$\frac{d\mu(t)}{dt} = \int_0^t \alpha_0 \theta(t-\tau) i(\tau) \mu(\tau) d\tau + \beta_0 e(t) \mu(t) \text{ and from Theorem 1(i) that}$$

$$\frac{d\lambda_0}{dt} = -w_s - \alpha_0 i(t) x_0(t) - \lambda_0(t) \beta_0 e(t). \text{ Substituting these expressions for } d\mu/dt \text{ and } d\lambda_1/dt \text{ into}$$

the expression for the sum of the first two terms of Theorem 3(i) given by

$$\beta_0 \left[ \lambda_0(t) \frac{d\mu(t)}{dt} + \frac{d\lambda_0(t)}{dt} \mu(t) \right] \text{ and eliminating common terms we obtain:}$$

$$\beta_0 \left[ \lambda_0(t) \left[ \int_0^t \alpha_0 \theta(t-\tau) i(\tau) \mu(\tau) d\tau \right] + \left[ -w_s - \alpha_0 i(t) \int_t^T \theta(\tau-t) \lambda_0(\tau) d\tau \right] \mu(t) \right]. \quad (\text{B18})$$

Applying the reverse time transformation to the fourth terms of Equation (B1) we obtain:

$$\lambda_0(t) \int_0^t \alpha_0 \theta(t-\tau) i(\tau) \mu(\tau) d\tau \equiv \alpha_0 i(t) \mu(t) \int_t^T \theta(\tau-t) \lambda_0(\tau) d\tau. \quad (\text{B19})$$

Since we know  $\beta_0, w_s, \mu(t) > 0$ , then substituting Equation (B19) into Equation (B18) and

$$\text{eliminating common terms we obtain: } \beta_0 \left[ \lambda_0(t) \frac{d\mu(t)}{dt} + \frac{d\lambda_0(t)}{dt} \mu(t) \right] \equiv -\beta_0 w_s \mu(t) < 0.$$

Therefore we have shown that the sum of the first two terms from Theorem 3(i) is negative. Given  $\beta_1, \sigma(t) > 0$ , we know that  $de(t)/dt < 0$  at time  $t \in [0, T]$  holds if the sum of the third and fourth terms for the variance effects of exploitation is negative. Clearly, if both the third and fourth terms are negative then their sum is negative. Therefore, it follows that  $de(t)/dt < 0$  at time  $t \in [0, T]$  if the conditions given in the statement of Corollary 4a hold.

**PROOF OF COROLLARY 4b.** The proof of Corollary 4b is analogous to the proof for Corollary 4a.

## Numerical Analysis

Numerical solutions are obtained using an ordinary shooting method for a discrete approximation of the continuous model (Sethi and Thompson 2000). We present a subset of the numerical analysis which highlights six temporal ambidexterity strategies. The parameter settings, short-term and long-term (terminal time) objectives for each scenario are presented in Table B2. The optimal rates of exploration and exploitation, ( $i(t)$ ,  $e(t)$ ), the instantaneous and cumulative distributed marginal values of the variance, ( $\lambda_1(t)$ ,  $x_1(t)$ ), and the evolution of the mean performance and the variance ( $\mu(t)$ ,  $\sigma(t)$ ) over the planning horizon process are illustrated in Figures B1-6 below. Note that all figures are drawn over time, for 20 time periods (i.e., the x-axis is  $t$ ).

The probability density function for the distributed lag  $\theta$ , is given by Equation (B20) below, where  $\omega$  and  $\kappa$  are the Gamma distribution's shape and scale parameters, respectively. The mean and variance of the probability density function are given by Equations (B21) and (B22) respectively. For a given shape parameter, as the scale factor increases, both the mean and the variance are larger. Therefore, an increasing scale factor represents a lagged effect in which a larger portion of the impact from investments in knowledge creation is realized later in the planning horizon.


$$\theta(x) = (x^{\omega-1} e^{-x/\kappa}) / \Gamma(\omega) \kappa^{\omega} \quad (B20)$$

$$\text{Mean}(X) = \omega \kappa \quad (B21)$$


$$\text{Variance}(X) = \omega \kappa^2 \quad (B22)$$

**Table B2: Numerical Analysis Parameter Settings**


Case	Short Term Objective	Long Term Objective	Exploration Parameters	Exploitation Parameters	Exploration Distributed Lag $\Theta \sim \text{Gamma}(\omega, \kappa)$	Initial Knowledge
1A	$P_S = -1$ $w_S = 0.002$	$P_L = -1$ $w_L = 0.002$	$\alpha_0 = 1$ $\alpha_1 = 5$ $c_0 = 20$	$\beta_0 = 2$ $\beta_1 = 8$ $c_1 = 20$	$\omega = 3$ $\kappa = 1$	$\mu(0) = 50$ $\sigma(0) = 45$
1B	$P_S = -1$ $w_S = 0.002$	$P_L = -1$ $w_L = 0.002$	$\alpha_0 = 3$ $\alpha_1 = 5$ $c_0 = 20$	$\beta_0 = 2$ $\beta_1 = 8$ $c_1 = 20$	$\omega = 3$ $\kappa = 1$	$\mu(0) = 50$ $\sigma(0) = 45$
1C	$P_S = -1$ $w_S = 0.002$	$P_L = -1$ $w_L = 0.002$	$\alpha_0 = 2$ $\alpha_1 = 5$ $c_0 = 20$	$\beta_0 = 2$ $\beta_1 = 8$ $c_1 = 20$	$\omega = 3$ $\kappa = 1$	$\mu(0) = 50$ $\sigma(0) = 5$
1D	$P_S = -1$ $w_S = 0.002$	$P_L = -1$ $w_L = 0.002$	$\alpha_0 = 2$ $\alpha_1 = 5$ $c_0 = 12$	$\beta_0 = 2$ $\beta_1 = 8$ $c_1 = 30$	$\omega = 3$ $\kappa = 1$	$\mu(0) = 50$ $\sigma(0) = 5$
1E	$P_S = -200$ $w_S = 0.00001$	$P_L = -1$ $w_L = 0.002$	$\alpha_0 = 3$ $\alpha_1 = 2$ $c_0 = 16$	$\beta_0 = 3$ $\beta_1 = 3$ $c_1 = 20$	$\omega = 3$ $\kappa = 0.01, 0.25 \text{ or } 1$	$\mu(0) = 50$ $\sigma(0) = 5$
2A	$P_S = 1$ $w_S = 0.002$	$P_L = 1$ $w_L = 0.002$	$\alpha_0 = 0.01$ $\alpha_1 = 3$ $c_0 = 20$	$\beta_0 = 4$ $\beta_1 = 2$ $c_1 = 20$	$\omega = 3$ $\kappa = 0.01, 0.25 \text{ or } 1$	$\mu(0) = 15$ $\sigma(0) = 15$
2B	$P_S = 1$ $w_S = 0.002$	$P_L = 1$ $w_L = 0.002$	$\alpha_0 = 0.01$ $\alpha_1 = 3$ $c_0 = 20$	$\beta_0 = 4$ $\beta_1 = 2$ $c_1 = 20$	$\omega = 3$ $\kappa = 1$	$\mu(0) = 15$ $\sigma(0) = 5$
3	$P_S = 10$ $w_S = 0.002$	$P_L = -10$ $w_L = 0.001$	$\alpha_0 = 3$ $\alpha_1 = 1$ $c_0 = 20$	$\beta_0 = 3$ $\beta_1 = 1$ $c_1 = 20$	$\omega = 3$ $\kappa = 1$	$\mu(0) = 40$ $\sigma(0) = 20$
4	$P_S = -50$ $w_S = 0.001$	$P_L = 10$ $w_L = 0.001$	$\alpha_0 = 3$ $\alpha_1 = 1$ $c_0 = 20$	$\beta_0 = 3$ $\beta_1 = 1$ $c_1 = 20$	$\omega = 3$ $\kappa = 1$	$\mu(0) = 60$ $\sigma(0) = 20$



*Instantaneous Marginal Variance,  $\lambda_I(t)$*




*Cumulative Marginal Variance,  $x_I(t)$*




*Rates of Exploration and Exploitation*

*Exploration,  $i(t)$  -●-    Exploitation,  $e(t)$  -□-*

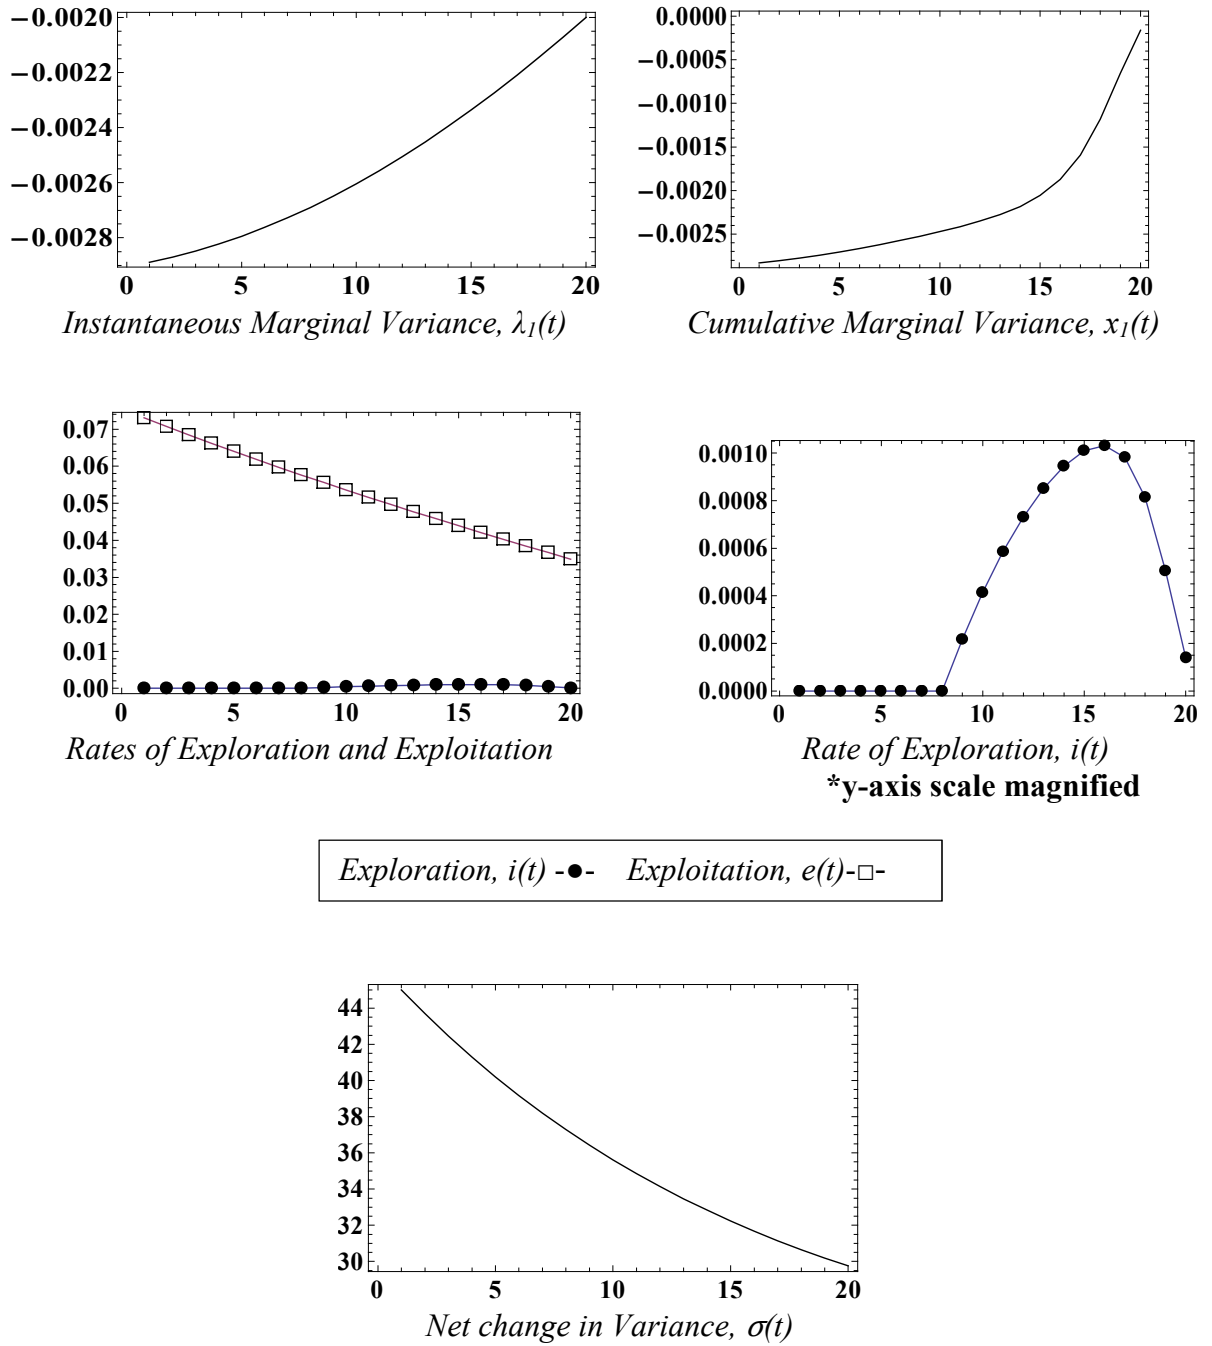


*Net change in Variance,  $\sigma(t)$*




*Lower Performance Bound,  $\mu(t)-\sigma(t)$ ;  
Mean,  $\mu(t)$ ;  
Upper Performance Bound,  $\mu(t)+\sigma(t)$ ;*


**Figure B5: Case 1A Short-Term Risk Averse and Long-Term Risk Averse**  
SCENARIO: Small Marginal Improvements in Mean from Exploration



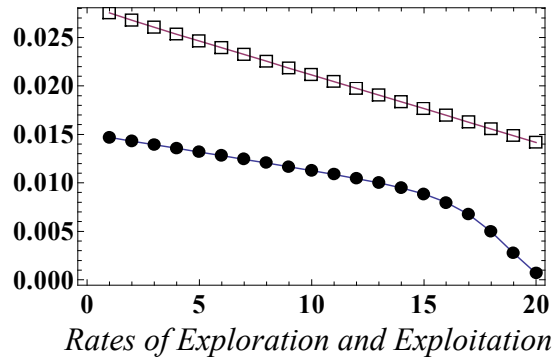
**Figure B6: Case 1B Short-Term Risk Averse and Long-Term Risk Averse**  
**SCENARIO: Large Initial Variance**



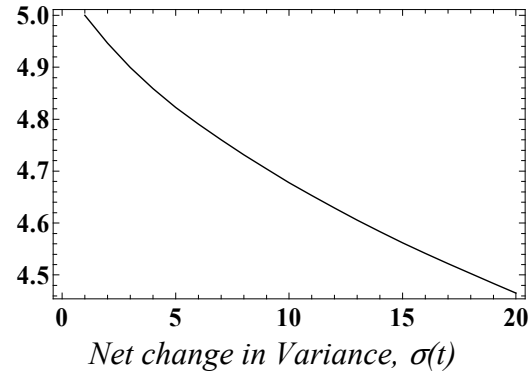
*Instantaneous Marginal Variance,  $\lambda_I(t)$*





*Cumulative Marginal Variance,  $x_I(t)$*




*Exploration,  $i(t)$  -●-      Exploitation,  $e(t)$ -□-*



**Figure B7: Case 1C Short-Term Risk Averse and Long-Term Risk Averse**  
SCENARIO: Small Initial Variance

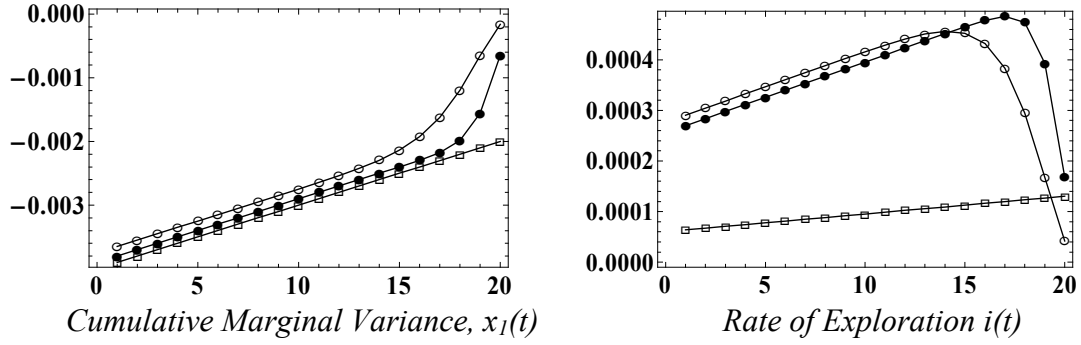


Exploration,  $i(t)$  -●-      Exploitation,  $e(t)$  -□-



**Figure B8: Case 1D Short-Term Risk Averse and Long-Term Risk Averse**  
 SCENARIO: Small Initial Variance; Low Cost of Exploration






No Lag -□-  $\omega = 3; \kappa = 0.01$

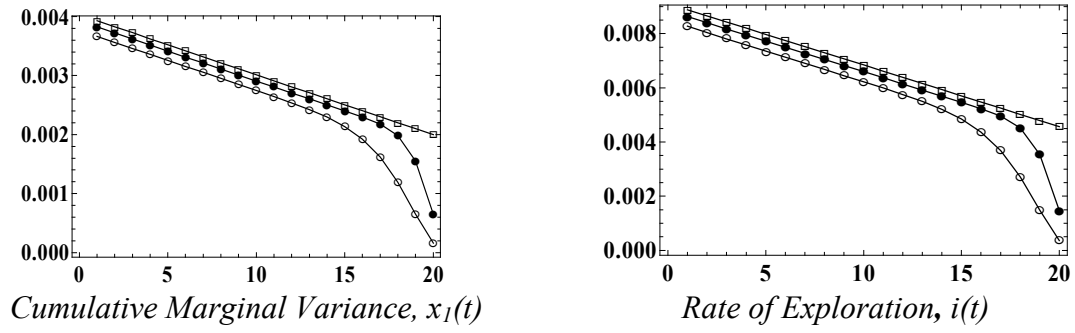
Small Lag -●-  $\omega = 3; \kappa = 0.25$

Large Lag -○-  $\omega = 3; \kappa = 1$



Exploration,  $i(t)$  -●-      Exploitation,  $e(t)$  -□-


**Figure B9: Case 1E Short-Term Risk Averse and Long-Term Risk Averse**  
SCENARIO: Effect of Exploration Lag



*No Lag* -□-  $\omega = 3; \kappa = 0.01$

*Small Lag* -●-  $\omega = 3; \kappa = 0.25$

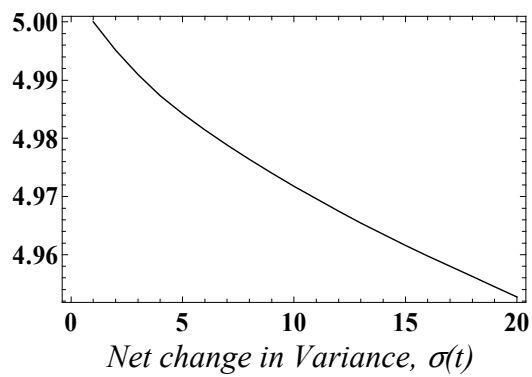

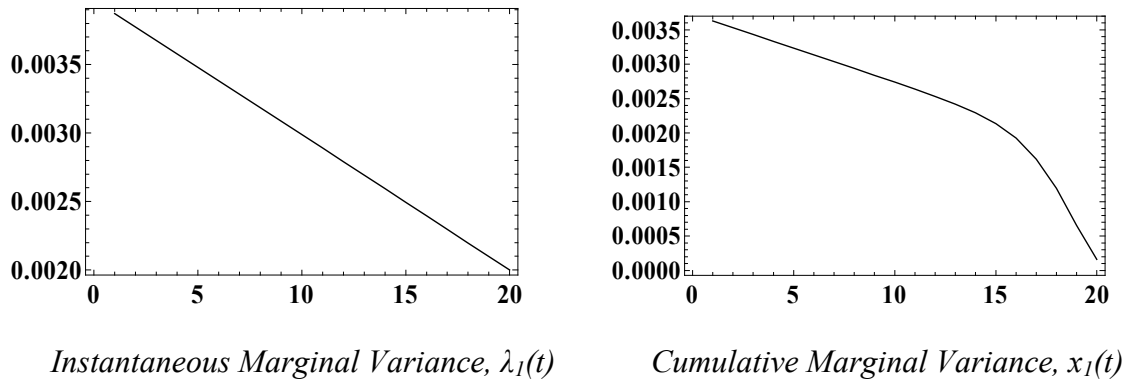
*Large Lag* -○-  $\omega = 3; \kappa = 1$




*Rates of Exploration  $i(t)$  and Exploitation  $e(t)$   
with increasing Gamma Distribution Scale Factor  $\kappa$*

*Exploration,  $i(t)$*  -●-    *Exploitation,  $e(t)$*  -□-


**Figure B10: Case 2A Short-Term Risk Seeking and Long-Term Risk Seeking  
SCENARIO: Effect of Exploration Lag**




**Figure B11: Case 2B Short-Term Risk Seeking and Long-Term Risk Seeking**  
SCENARIO: Small Variance



*Instantaneous Marginal Variance,  $\lambda_1(t)$*




*Cumulative Marginal Variance,  $x_1(t)$*





*Rates of Exploration and Exploitation*


<i>Exploration, <math>i(t)</math> -●-</i>	<i>Exploitation, <math>e(t)</math> -□-</i>
---	--



**Figure B12: Case 3 Short Term Risk Seeking and Long Term Risk Averse**



Exploration,  $i(t)$  -●-      Exploitation,  $e(t)$  -□-



**Figure B13: Case 4 Short Term Risk Averse and Long Term Risk Seeking**

## APPENDIX C

**Table C1: Model Notation**

$t$	Time, $t \in [0, T]$ : 0 (T) denotes the start (end) of the planning horizon.
<b>Parameters for Focal Firm, <math>j=1</math></b>	
$i(t)$	Rate of exploration efforts at time $t$ , $i(t) \geq 0$ ; control variable.
$e(t)$	Rate of exploitation efforts at time $t$ , $e(t) \geq 0$ ; control variable.
$\mu(t)$	Mean technical performance at time $t$ , $\mu(t) \geq 0$ ; $\mu(0)$ given; state variable.
$\sigma(t)$	Variance technical performance at time $t$ , $\sigma(t) \geq 0$ ; $\sigma(0)$ given; state variable.
$v(t)$	Random variable indicating the innovation's technical performance at time $t$ ; $v(t) \sim N(\mu(t), \sigma(t))$ .
$\alpha_0$ ( $\alpha_1$ )	Marginal impact of exploration on the mean (variance); $\alpha_0$ ( $\alpha_1$ ) $> 0$ .
$\beta_0$ ( $\beta_1$ )	Marginal impact of exploitation on the mean (variance); $\beta_0$ ( $\beta_1$ ) $> 0$ .
$\psi$	Degree of participation in exploration knowledge-sharing alliance
$\gamma$	Degree of participation in exploitation knowledge-sharing alliance
$c_0$	Parameter indicating marginal cost of exploration.
$c_1$	Parameter indicating marginal cost of exploitation.
$\lambda_\mu(t)$	Marginal value to the focal firm of a unit increase in the mean of its technical performance at time $t$ .
$\lambda_\sigma(t)$	Marginal value to the focal firm of a unit increase in the variance of its technical performance at time $t$ .
$\lambda_M(t)$	Marginal value to the focal firm of a unit increase in the mean of the rival's technical performance at time $t$ .
$\lambda_S(t)$	Marginal value to the focal firm of a unit increase in the variance of the rival's technical performance at time $t$ .
<b>Parameters for Rival Firm, <math>j=2</math></b>	
$I(t)$	Rate of exploration efforts at time $t$ , $I(t) \geq 0$ ; control variable.
$E(t)$	Rate of exploitation efforts at time $t$ , $E(t) \geq 0$ ; control variable.
$M(t)$	Mean technical performance at time $t$ , $M(t) \geq 0$ ; $M(0)$ given; state variable.
$S(t)$	Variance technical performance at time $t$ , $S(t) \geq 0$ ; $S(0)$ given; state variable.

$V(t)$	Random variable indicating the innovation's technical performance at time $t$ ; $V(t) \sim N(M(t), S(t))$ .
$a_0 (a_1)$	Marginal impact of exploration on the mean (variance); $a_0 (a_1) > 0$ .
$b_0 (b_1)$	Marginal impact of exploitation on the mean (variance); $b_0 (b_1) > 0$ .
$X$	Degree of participation in exploration knowledge-sharing alliance
$Y$	Degree of participation in exploitation knowledge-sharing alliance
$C_0$	Parameter indicating marginal cost of exploration.
$C_1$	Parameter indicating marginal cost of exploitation.
$L_M(t)$	Marginal value to the rival firm of a unit increase in the mean of its technical performance at time $t$ .
$L_S(t)$	Marginal value to the rival firm of a unit increase in the variance of its technical performance at time $t$ .
$L_\mu(t)$	Marginal value to the rival firm of a unit increase in the mean of the focal firm's technical performance at time $t$ .
$L_\sigma(t)$	Marginal value to the rival firm of a unit increase in the variance of the focal firm's technical performance at time $t$ .

### Hamiltonian

The Hamiltonian function,  $H$ , is given below for  $j=1$ . The Hamiltonian for  $j=2$  is analogous.

$$\begin{aligned} H = & -1/2c_0 i^2(t) - 1/2c_1 e^2(t) \\ & + \lambda_\mu(t)(\alpha_0 i(t) + \beta_0 e(t)) + \lambda_\sigma(t)(\alpha_1 i(t)S^x(t) - \beta_1 e(t)S^{-y}(t)) \\ & + \lambda_M(t)(a_0 I(t) + b_0 E(t)) + \lambda_s(t)(a_1 I(t)\sigma^\psi(t) - b_1 E(t)\sigma^{-\gamma}(t)) \end{aligned}$$

### Lagrangian

The Lagrangian function to be maximized is given below for  $j=1$ . The Lagrangian for  $j=2$  is analogous.

$$L = H + \eta_\sigma(t)[\alpha_1 i(t)S^x(t) - \beta_1 e(t)S^{-y}(t)] + \eta_s(t)[a_1 I(t)\sigma^\psi(t) - b_1 E(t)\sigma^{-\gamma}(t)]$$

### Optimality Conditions

The optimality conditions for the open loop dynamic game are as follows for  $j=1$ .

$$\frac{d\lambda_\mu(t)}{dt} = \frac{-dL}{d\mu(t)} \text{ and } \lambda_\mu(T) = \frac{df}{d\mu(T)} \quad (C1)$$

$$\frac{d\lambda_\sigma(t)}{dt} = \frac{-dL}{d\sigma(t)} \text{ and } \lambda_\sigma(T) = \frac{df}{d\sigma(T)} \quad (C2)$$

$$\frac{d\lambda_M(t)}{dt} = \frac{-dL}{dM(t)} \text{ and } \lambda_M(T) = \frac{df}{dM(T)} \quad (C3)$$

$$\frac{d\lambda_s(t)}{dt} = \frac{-dL}{dS(t)} \text{ and } \lambda_s(T) = \frac{df}{dS(T)} \quad (C4)$$

$$\partial L / \partial e(t) = 0 \quad (C5)$$

$$\partial L / \partial i(t) = 0 \quad (C6)$$

$$\eta_\sigma(t)\sigma(t) = 0; \quad \eta_s(t)S(t) = 0; \quad (C7)$$

$$\eta_\sigma(t) \geq 0; \quad \eta_s(t) \geq 0 \quad (C8)$$

$$d\eta_\sigma(t)/dt < 0; \quad \eta_s(t)/dt < 0 \quad (C9)$$

$$\sigma(t) > 0; \quad S(t) > 0 \quad (C10)$$

$$e(t) \geq 0, \quad i(t) \geq 0 \quad (C11)$$



Note that these non-negative constraints and optimality conditions ensure  $\sigma(t), S(t) > 0$  hold. However, for simplicity in the remainder of the paper, we assume  $\sigma(t), S(t) > 0$  hold. Therefore, we optimize the Hamiltonian and do not consider the Lagrangian (i.e.,  $\eta_\sigma(t), \eta_S(t) = 0$  for all  $t \in [0, T]$ ). Lastly, for the sufficiency condition given in Feichtinger and Jørgensen (1983) to hold in an exploration knowledge-sharing alliance where  $\gamma = 0$  and  $Y = 0$ , then the following condition must be satisfied for the focal firm  $j = 1$ :  $(X-1)\alpha_0\lambda_\mu(t) + (2X-1)\alpha_1\lambda_\sigma(t)S^X < 0$ . (Analogous conditions exist for firm  $j = 2$ , and for the exploitation knowledge-sharing alliance where  $\psi = 0$  and  $X = 0$ .) Although we are not able to analytically prove that the sufficiency conditions hold, in extensive numerical analysis in which each input parameter was varied over a wide range of values, there was not a single solution in which sufficiency was violated.

**Proof of Theorem 1.** The results of Theorem 1 follow from the optimality conditions in Equations (C1)-(C6) and (C11) given  $\eta_\sigma(t), \eta_S(t) = 0$ , for  $t \in [0, T]$ . Note that Theorem 1 focuses on the scenario where  $\gamma = 0, Y = 0$ .

**Proof of Corollary 1A.** We know that  $\gamma = 0, Y = 0$  and Equation (3) hold. From Equations (3), (8) and (9), we obtain:

$$\lambda_\sigma(T) > 0, d\lambda_\sigma(T)/dt > 0, \lambda_S(T) < 0 \text{ and } d\lambda_S(T)/dt < 0. \quad (\text{C12})$$

Also from Equations (8) and (9), we observe that  $\lambda_\sigma(t)$  and  $\lambda_S(t)$  must be solved simultaneously. This gives us the following for any  $t \in [0, T]$ :

$$\text{If } \lambda_\sigma(t) > 0 \text{ then } d\lambda_S(t)/dt < 0. \quad (\text{C13i})$$

$$\text{If } \lambda_S(t) > 0 \text{ then } d\lambda_\sigma(t)/dt < 0. \quad (\text{C13ii})$$

$$\text{If } \lambda_\sigma(t) < 0 \text{ then } d\lambda_S(t)/dt > 0. \quad (\text{C13iii})$$

$$\text{If } \lambda_S(t) < 0 \text{ then } d\lambda_\sigma(t)/dt > 0. \quad (\text{C13iv})$$

$$\text{If } \lambda_\sigma(t) = 0 \text{ then } d\lambda_S(t)/dt = 0. \quad (\text{C13v})$$


$$\text{If } \lambda_S(t) = 0 \text{ then } d\lambda_\sigma(t)/dt = 0. \quad (\text{C13vi})$$

Lastly, given Equation (C12) and the conditions in (C13i)-(C13vi) above, we obtain the following pairs of solutions for  $\lambda_\sigma(t)$  and  $\lambda_s(t)$ , for  $t \in [0, T]$  depicted in Figure C1:


Case 1A-I:  $\lambda_\sigma(t) > 0$ ,  $d\lambda_\sigma(t)/dt \geq 0$ ,  $\lambda_s(t) < 0$ ,  $d\lambda_s(t)/dt \leq 0$  for  $t \in [0, T]$ .

Case 1A-II:  $\lambda_\sigma(t) \leq 0$ ,  $d\lambda_\sigma(t)/dt \geq 0$ ,  $\lambda_s(t) < 0$ ,  $d\lambda_s(t)/dt \geq 0$  for  $t \in [0, t_1]$ ;  $\lambda_\sigma(t) > 0$ ,  $d\lambda_\sigma(t)/dt \geq 0$ ,  $\lambda_s(t) < 0$ ,  $d\lambda_s(t)/dt < 0$  for  $t \in (t_1, T]$ .


Case 1A-III:  $\lambda_\sigma(t) > 0$ ,  $d\lambda_\sigma(t)/dt \leq 0$ ,  $\lambda_s(t) \geq 0$ ,  $d\lambda_s(t)/dt \leq 0$  for  $t \in [0, t_1]$ ;  $\lambda_\sigma(t) > 0$ ,  $d\lambda_\sigma(t)/dt > 0$ ,  $\lambda_s(t) < 0$ ,  $d\lambda_s(t)/dt \leq 0$  for  $t \in (t_1, T]$ .



**Case 1A-I**




**Case 1A-II**



**Case 1A-III**

**Figure C1: Solutions for  $\lambda_\sigma(t)$  and  $\lambda_s(t)$  under Case 1A**

The only other solutions that satisfy the conditions for  $\lambda_\sigma(T)$  and  $d\lambda_\sigma(T)/dt$  given in Equation (C12) are illustrated in Figure C2 below.



**Figure C2: Possible solutions for  $\lambda_\sigma(t)$**

Consider the solution for  $\lambda_\sigma(t)$  shown in Figure C2. With  $\lambda_\sigma(t) < 0$  and  $d\lambda_\sigma(t)/dt < 0$  for  $t \in [0, t_3)$ , from Equations (C13ii) and (C13iii) we obtain  $\lambda_s(t) > 0$  and  $d\lambda_s(t)/dt > 0$  for  $t \in [0, t_3)$ . In addition, with  $\lambda_\sigma(t_3) < 0$  and  $d\lambda_\sigma(t_3) = 0$ , from Equations (C13iii) and (C13vi) we obtain  $\lambda_s(t_3) = 0$  and  $d\lambda_s(t_3)/dt > 0$ . However, since  $\lambda_s^*(t)$  is a continuous function over all  $t \in [0, T]$  these conditions cannot occur. Therefore the solution for  $\lambda_\sigma(t)$  illustrated in Figure C2-IV cannot occur. Similarly, it can be shown that the solution for  $\lambda_\sigma(t)$  illustrated in Figure C2-V cannot occur.

In conclusion, the only solutions of  $\lambda_\sigma(t)$  and  $d\lambda_s(t)$  for  $t \in [0, T]$ , corresponding to the conditions given in Corollary 1A can occur. QED

**Proof of Corollary 1B.** We know that  $\gamma = 0$ ,  $Y = 0$  and Equation (4) holds. From Equations (4), (8) and (9), we obtain:


$$\lambda_\sigma(T) < 0, d\lambda_\sigma(T)/dt \leq 0, \lambda_s(T) > 0 \text{ and } d\lambda_s(T)/dt \geq 0 \quad (\text{C14})$$

Since  $\lambda_\sigma(t)$  and  $\lambda_s(t)$  must hold simultaneously, then by analogous reasoning to Corollary 1A, there are three possible solutions for both  $\lambda_\sigma(t)$  and  $\lambda_s(t)$  which are continuous functions for  $t \in [0, T]$ :


Case 1B-I:  $\lambda_\sigma(t) < 0$ ,  $d\lambda_\sigma(t)/dt \leq 0$ ,  $\lambda_s(t) > 0$ ,  $d\lambda_s(t)/dt \geq 0$  for  $t \in [0, T]$ .

Case 1B-II:  $\lambda_\sigma(t) < 0$ ,  $d\lambda_\sigma(t)/dt \geq 0$ ,  $\lambda_s(t) \leq 0$ ,  $d\lambda_s(t)/dt \geq 0$  for  $t \in [0, t_1]$  and  $\lambda_\sigma(t) < 0$ ,  $d\lambda_\sigma(t)/dt < 0$ ,  $\lambda_s(t) > 0$ ,  $d\lambda_s(t)/dt \geq 0$  for  $t \in (t_1, T]$ .


Case 1B-III:  $\lambda_\sigma(t) \geq 0$ ,  $d\lambda_\sigma(t)/dt \leq 0$ ,  $\lambda_s(t) > 0$ ,  $d\lambda_s(t)/dt \leq 0$  for  $t \in [0, t_2]$  and  $\lambda_\sigma(t) < 0$ ,  $d\lambda_\sigma(t)/dt \leq 0$ ,  $\lambda_s(t) > 0$ ,  $d\lambda_s(t)/dt > 0$  for  $t \in (t_2, T]$ .



**Case 1B-I**



**Case 1B-II**



**Case 1B-III**

**Figure C3: Solutions for  $\lambda_\sigma(t)$  and  $\lambda_s(t)$  under Case 1B**

**Proof of Theorem 2.** The results of Theorem 2 follow from the optimality conditions in Equations (C5) and (C6) and non-negativity constraints on  $i_j(t)$  and  $e_j(t)$  and with  $\psi=0$ ,  $X=0$ .

**Proof of Corollary 2A.** We know that  $\psi=0$ ,  $X=0$  and Equation (3) holds. From Equations (3), (12) and (13), we obtain the following:

$$\lambda_\sigma(T)>0, d\lambda_\sigma(T)/dt\leq 0, \lambda_S(T)<0 \text{ and } d\lambda_S(T)/dt\geq 0. \quad (C15)$$

Since  $\lambda_\sigma(t)$  and  $\lambda_S(t)$  must hold simultaneously, then by analogous reasoning to Corollary 1A, there are three possible solutions for both  $\lambda_\sigma(t)$  and  $\lambda_S(t)$  which are continuous functions for  $t\in[0,T]$ :

Case 2A-I:  $\lambda_\sigma(t)>0, d\lambda_\sigma(t)/dt\geq 0, \lambda_S(t)<0, d\lambda_S(t)/dt\leq 0$  for  $t\in[0,T]$ .

Case 2A-II:  $\lambda_\sigma(t)\leq 0, d\lambda_\sigma(t)/dt\geq 0, \lambda_S(t)<0, d\lambda_S(t)/dt\geq 0$  for  $t\in[0,t_1]$ ;  $\lambda_\sigma(t)>0, d\lambda_\sigma(t)/dt\geq 0, \lambda_S(t)<0, d\lambda_S(t)/dt<0$  for  $t\in(t_1,T]$ .

Case 2A-III:  $\lambda_\sigma(t)>0, d\lambda_\sigma(t)/dt\leq 0, \lambda_S(t)\geq 0, d\lambda_S(t)/dt\leq 0$  for  $t\in[0,t_1]$ ;  $\lambda_\sigma(t)>0, d\lambda_\sigma(t)/dt>0, \lambda_S(t)<0, d\lambda_S(t)/dt\leq 0$  for  $t\in(t_1,T]$ .

**Proof of Corollary 2B.** We know that  $\psi=0$ ,  $X=0$  and Equation (4) holds. From Equations (4), (12) and (13), we obtain the following:

$$\lambda_\sigma(T)<0, d\lambda_\sigma(T)/dt\leq 0, \lambda_S(T)>0 \text{ and } d\lambda_S(T)/dt\geq 0 \quad (C16)$$

Since  $\lambda_\sigma(t)$  and  $\lambda_S(t)$  must hold simultaneously, then by analogous reasoning to Corollary 1A, there are three possible solutions for both  $\lambda_\sigma(t)$  and  $\lambda_S(t)$  which are continuous functions for  $t\in[0,T]$ :

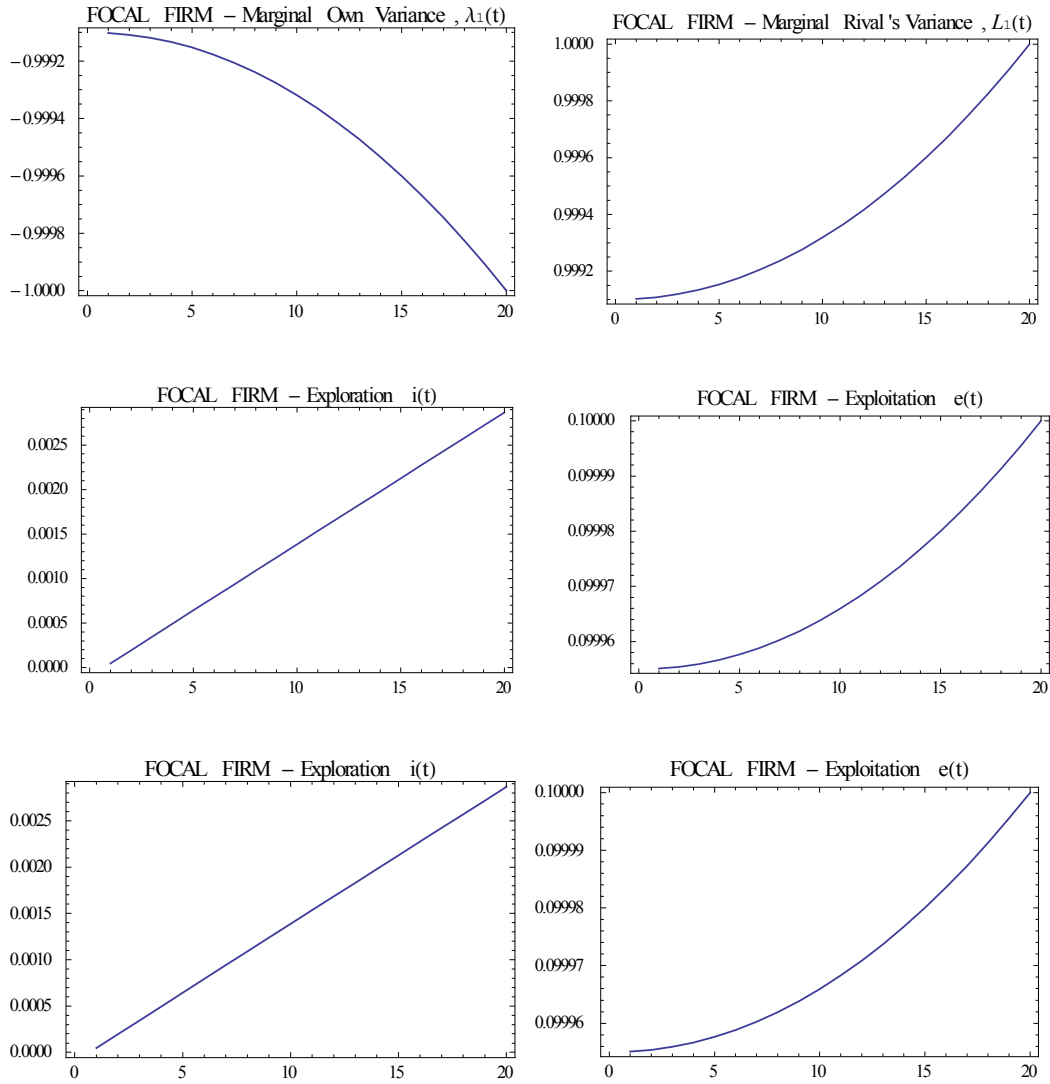
Case 2B-I:  $\lambda_\sigma(t)<0, d\lambda_\sigma(t)/dt\leq 0, \lambda_S(t)>0, d\lambda_S(t)/dt\geq 0$  for  $t\in[0,T]$ .

Case 2B-II:  $\lambda_\sigma(t)<0, d\lambda_\sigma(t)/dt\geq 0, \lambda_S(t)\leq 0, d\lambda_S(t)/dt\geq 0$  for  $t\in[0,t_1]$  and  $\lambda_\sigma(t)<0, d\lambda_\sigma(t)/dt<0, \lambda_S(t)>0, d\lambda_S(t)/dt\geq 0$  for  $t\in(t_1,T]$ .

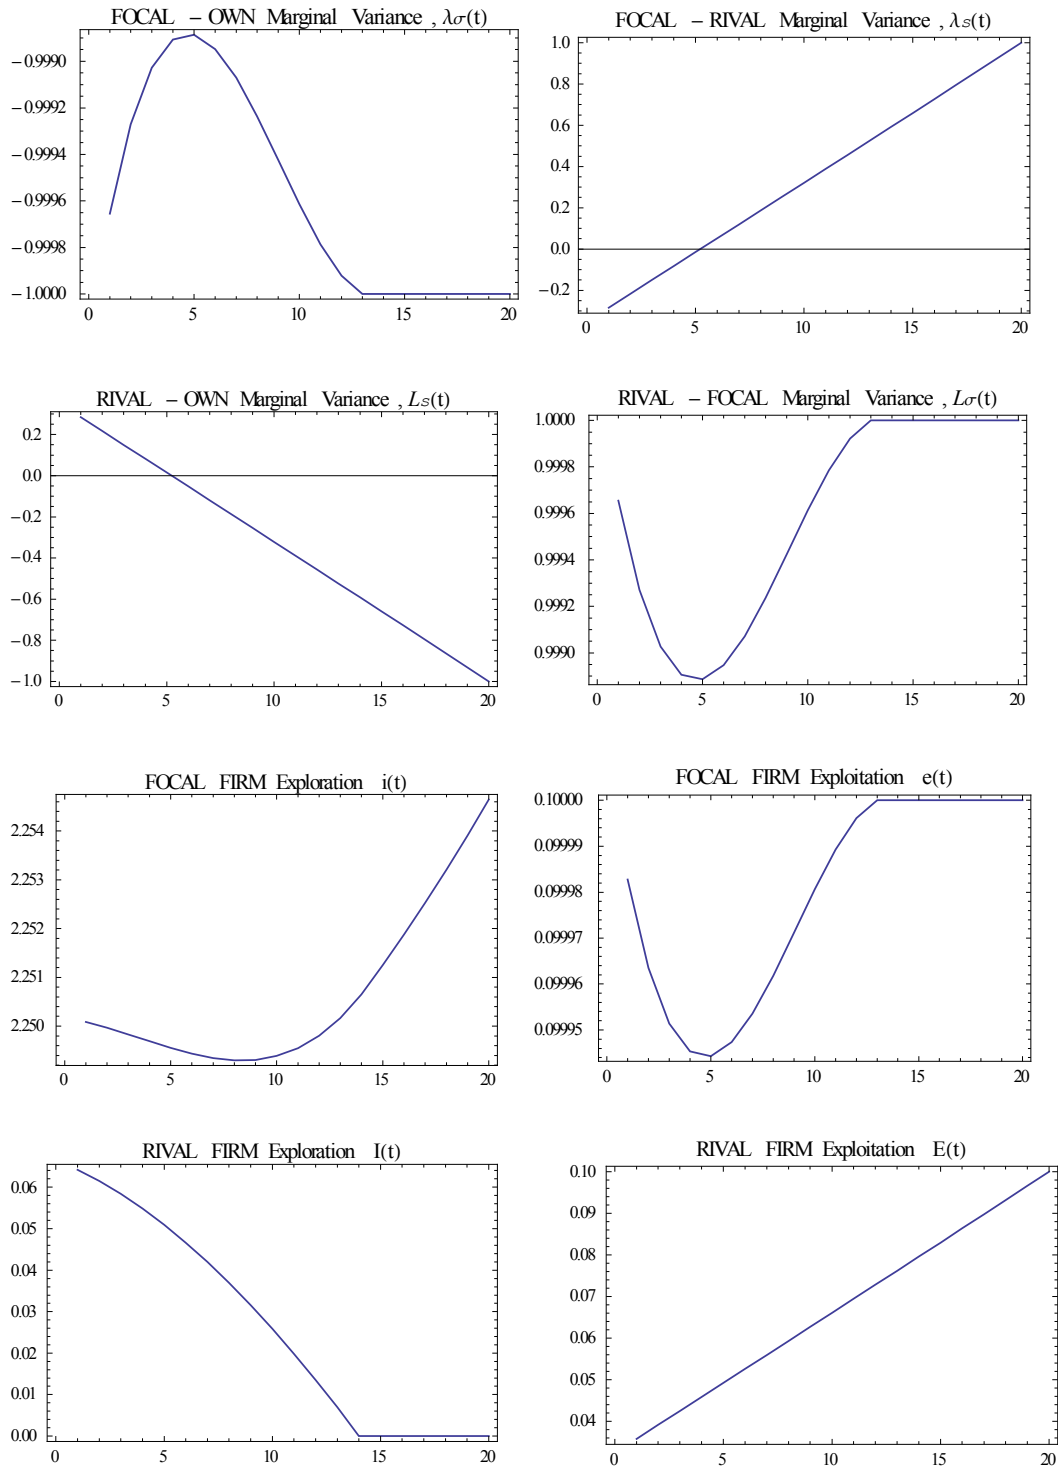
Case 2B-III:  $\lambda_\sigma(t)\geq 0, d\lambda_\sigma(t)/dt\leq 0, \lambda_S(t)>0, d\lambda_S(t)/dt\leq 0$  for  $t\in[0,t_2]$  and  $\lambda_\sigma(t)<0, d\lambda_\sigma(t)/dt\leq 0, \lambda_S(t)>0, d\lambda_S(t)/dt>0$  for  $t\in(t_2,T]$ .

**Table C2: Numerical Analysis Parameter Settings**


	FOCAL FIRM PARAMETERS										RIVAL FIRM PARAMETERS									
	$\alpha_0$	$\alpha_1$	$c_0$	$\beta_0$	$\beta_1$	$c_1$	$\mu(0)$	$\sigma(0)$	$\psi$	$\gamma$	$a_0$	$a_1$	$C_0$	$b_0$	$b_1$	$C_1$	$M(0)$	$S(0)$	$X$	$Y$
1B-I	1	1	20	1	1	20	1	1	0.6	0.6	1	1	20	1	1	20	1	1	0.6	0.6
1B-II/III	10	1	4	1	1	20	1	1	0.6	0.6	1	1	20	1	1	20	1	1	0.6	0.6



**Figure C4: Case 1B-I**



**Figure C5: Case 1B-II for firm j=1 and Case 1B-III for firm j=2**



**Figure C5 (continued): Case 1B-II for firm  $j=1$  and Case 1B-III for firm  $j=2$**



## REFERENCES

- Abbott Laboratories. Abbott Completes Acquisition of IDEV Technologies [Press release]. August 21, 2013. Retrieved from <http://www.abbott.com/press-release/abbott-completes-acquisition-of-idev-technologies.htm>
- Abdellaoui, M., Diecidue, E., and Öncüler, A. 2011. Risk preferences at different time periods: An experimental investigation. *Management Science*, 57(5): 975-987.
- Adler, P. S., and Clark, K. B. 1991. Behind the learning curve: A sketch of the learning process. *Management Science*, 37(3): 267-281.
- Ahuja, G., and Lampert, C.M. 2001. Entrepreneurship in the large corporation: A longitudinal study of how established firms create breakthrough inventions. *Strategic Management Journal*, 22(6-7): 521-543.
- Aiken, L. S., and West, S. G. 1991. Multiple regression: Testing and interpreting interactions: Sage Publications, Incorporated.
- Albert, M. B., Avery, D., Narin, F., and McAllister, P. 1991. Direct validation of citation counts as indicators of industrially important patents. *Research Policy*, 20(3): 251-259.
- Alexy, O., George, G., and Salter, A. J. 2013. Cui bono? The selective revealing of knowledge and its implications for innovative activity. *Academy of Management Review*, 38(2): 270-291.
- Ali, A. 1994. Pioneering versus incremental innovation: review and research propositions. *Journal of Product Innovation Management*, 11(1): 46-61.
- Ali, A., Kalwani, M. U., and Kovenock, D. 1993. Selecting product development projects: Pioneering versus incremental innovation strategies. *Management Science*, 39(3): 255-274.
- Anderson, P., and Tushman, M. L. 1990. Technological discontinuities and dominant designs: A cyclical model of technological change. *Administrative Science Quarterly*, 31(3):439-465
- Andriopoulos, C., and Lewis, M. W. 2009. Exploitation-exploration tensions and organizational ambidexterity: Managing paradoxes of innovation. *Organization Science*, 20(4): 696-717.
- Argote, L. 2013. Organizational learning: Creating, retaining and transferring knowledge: Springer.

- Artz, K. W., Norman, P. M., Hatfield, D. E., and Cardinal, L. B. 2010. A longitudinal study of the impact of R&D, patents, and product innovation on firm performance. *Journal of Product Innovation Management*, 27(5): 725-740.
- Audia, P. G., and Goncalo, J. A. 2007. Past success and creativity over time: A study of inventors in the hard disk drive industry. *Management Science*, 53(1): 1-15.
- Autio, E., Hameri, A. P., and Vuola, O. 2004. A framework of industrial knowledge spillovers in big-science centers. *Research Policy*, 33(1): 107-126.
- Austin, R. D., Devin, L., and Sullivan, E. E. 2012. Accidental innovation: Supporting valuable unpredictability in the creative process. *Organization Science*, 23(5): 1505-1522.
- Azoulay, P., Graff Zivin, J. S., and Manso, G. 2011. Incentives and creativity: evidence from the academic life sciences. *The RAND Journal of Economics*, 42(3): 527-554.
- Bajaj, A., Kekre, S., and Srinivasan, K. 2004. Managing NPD: Cost and schedule performance in design and manufacturing. *Management Science*, 50(4): 527-536.
- Barney, J. 1991. Firm resources and sustained competitive advantage. *Journal of Management*, 17(1): 99-120.
- Baron, R. A., and Ensley, M. D. 2006. Opportunity recognition as the detection of meaningful patterns: Evidence from comparisons of novice and experienced entrepreneurs. *Management Science*, 52(9): 1331-1344.
- Baron, R. M., and Kenny, D. A. 1986. The moderator–mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *Journal of Personality and Social Psychology*, 51(6): 1173.
- Baum, J. A. C., Calabrese, T., and Silverman, B. S. 2000. Don't go it alone: Alliance network composition and startups' performance in Canadian biotechnology. *Strategic Management Journal*, 21(3): 267-294.
- Baum, J. A., and Ingram, P. 1998. Survival-enhancing learning in the Manhattan hotel industry, 1898–1980. *Management Science*, 44(7), 996-1016.
- Baumard, P., and Starbuck, W. H. 2005. Learning from failures: Why it may not happen. *Long Range Planning*, 38(3), 281-298.
- Baumol, W. J. 1963. An expected gain-confidence limit criterion for portfolio selection. *Management Science* 10(1): 174-182.

- Benner, M. J., and Tushman, M. L. 2003. Exploitation, exploration, and process management: The productivity dilemma revisited. *Academy of Management Review*, 28(2): 238-256.
- Bessen, J. 2009. User documentation: Matching patent data to Compustat firms. Retrieved from <http://users.nber.org/~jbessen/matchdoc.pdf>.
- Biazzo, S. 2009. Flexibility, structuration, and simultaneity in new product development. *Journal of Product Innovation Management*, 26(3): 336-353.
- Bingham, C. B., and Davis, J. P. 2012. Learning sequences: their existence, effect, and evolution. *Academy of Management Journal*, 55(3): 611-641.
- Bohn, R. E. 1994. Measuring and managing technological knowledge. *Sloan Management Review*, 36(1): 61-73.
- Boudreau, K. J., Lacetera, N., and Lakhani, K. R. 2011. Incentives and problem uncertainty in innovation contests: An empirical analysis. *Management Science*, 57(5): 843-863.
- Boyle, E., and Shapira, Z. 2012. The liability of leading: Battling aspiration and survival goals in the Jeopardy! Tournament of Champions. *Organization Science*, 23(4): 1100-1113.
- Brockner, J., Higgins, E. T., and Low, M. B. 2004. Regulatory focus theory and the entrepreneurial process. *Journal of Business Venturing*, 19(2): 203-220.
- 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): 1-34.
- Browning, T. R., Deyst, J. J., Eppinger, S. D., and Whitney, D. E. 2002. Adding value in product development by creating information and reducing risk. *Engineering Management, IEEE Transactions*, 49(4): 443-458.
- Businesswire. Medtronic Signs Agreement to Acquire Ardian [Press release]. November 22, 2010. Retrieved from <http://www.businesswire.com>.
- Cabral, L. 2003. R&D competition when firms choose variance. *Journal of Economics and Management Strategy*, 12(1): 139-150.
- Cannon, M. D., and Edmondson, A. C. 2001. Confronting failure: Antecedents and consequences of shared beliefs about failure in organizational work groups. *Journal of Organizational Behavior*, 22(2): 161-177.

- Cannon, M. D., and Edmondson, A. C. 2005. Failing to learn and learning to fail (intelligently): How great organizations put failure to work to innovate and improve. *Long Range Planning*, 38(3): 299-319.
- Cao, Q., Gedajlovic, E., and Zhang, H. 2009. Unpacking organizational ambidexterity: Dimensions, contingencies, and synergistic effects. *Organization Science*, 20(4): 781-796.
- Carrillo, J. E., and Gaimon, C. 2004. Managing knowledge-based resource capabilities under uncertainty. *Management Science* 50(11): 1504-1518.
- Cellini, R., and Lambertini, L. 2002. A differential game approach to investment in product differentiation. *Journal of Economic Dynamics and Control*, 27(1): 51-62.
- Chao, R. O., Kavadias, S., and Gaimon, C. 2009. Revenue driven resource allocation: Funding authority, incentives, and new product development portfolio management. *Management Science*, 55(9): 1556-1569.
- Chandrasekaran, A., Linderman, K., and Schroeder, R. 2012. Antecedents to ambidexterity competency in high technology organizations. *Journal of Operations Management*, 30(1): 134-151.
- Choi, Y. R., Lévesque, M., Shepherd, D. A. 2008. When should entrepreneurs expedite or delay opportunity exploitation? *Journal of Business Venturing*, 23(3): 333-355.
- Cohen, W. M., Levinthal, D. A. 1990. Absorptive capacity: a new perspective on learning and innovation. *Administrative Science Quarterly*, 35(1). 128-152.
- Cohen, M. D., and Sproull, L.S. 1995. Organizational learning. Sage Publications. Thousand Oaks, CA.
- Colombo, M. G., Grilli, L., and Piva, E. 2006. In search of complementary assets: The determinants of alliance formation of high-tech start-ups. *Research Policy*, 35(8): 1166-1199.
- Cooper, R. G. 1990. Stage-gate systems: a new tool for managing new products. *Business Horizons*, 33(3): 44-54.
- Cross, N. 2000. Engineering design methods: strategies for product design. Wiley Chichester.
- Das, T., and Teng, B. S. 1997. Time and entrepreneurial risk behavior. *Entrepreneurship Theory and Practice*, 22: 69-88.

- Deeds, D. L., and Hill, C. W. 1996. Strategic alliances and the rate of new product development: an empirical study of entrepreneurial biotechnology firms. *Journal of Business Venturing*, 11(1), 41-55.
- Denend L., and Zenios S. 2006. Drug eluting stents: A paradigm shift in the medical device industry. Stanford Graduate School of Business, Case Study OIT-50, 02/13/0
- Dickson, P. R., and Giglierano, J. J. 1986. Missing the boat and sinking the boat: A conceptual model of entrepreneurial risk. *The Journal of Marketing*: 50(3): 58-70.
- Dosi, G.1982. Technological paradigms and technological trajectories: A suggested interpretation of the determinants and directions of technical change. *Research Policy*, 11(3): 147-162.
- Dunlap-Hinkler, D., Kotabe, M., and Mudambi, R. 2010. A story of breakthrough versus incremental innovation: Corporate entrepreneurship in the global pharmaceutical industry. *Strategic Entrepreneurship Journal*, 4(2): 106-127.
- Eesley, C., E. Roberts. 2010. Cutting your teeth: Learning from entrepreneurial experiences. *Academy of Management Annual Meeting Proceedings*, 2010, 1–6.
- Erat, S. and Kavadias, S. 2008. Sequential testing of product designs: Implications for learning. *Management Science*, 54(5): 956.
- Erat, S., and Krishnan, V. 2012. Managing delegated search over design spaces. *Management Science*, 58(3): 606-623.
- Feichtinger, G., and Jørgensen, S. 1983. Differential game models in management science. *European Journal of Operational Research*, 14(2), 137-155.
- Fixson, S. K., and Marion, T. J. 2012. Back-loading: A potential side effect of employing digital design tools in new product development. *Journal of Product Innovation Management*, 29(S1): 140-156.
- Fleming, L. 2001. Recombinant uncertainty in technological search. *Management Science*, 47(1): 117-132.
- Fleming, L., and Sorenson, O. 2001. Technology as a complex adaptive system: evidence from patent data. *Research Policy*, 30(7): 1019-1039.
- Fleming, L., and Sorenson, O. 2004. Science as a map in technological search. *Strategic Management Journal*, 25(8-9): 909-928.
- Forlani, D., Mullins, J. 2000. Perceived risks and choices in entrepreneurs' new venture decisions. *Journal of Business Venturing*, 15(4): 305-322.

- Franco, A. M., Sarkar, M., Agarwal, R., and Echambadi, R. 2010. Swift and smart: the moderating effects of technological capabilities on the market pioneering-firm survival relationship. *Management Science*, 55(11): 1842.
- Fried, Y., and Slowik, L. H. 2004. Enriching goal-setting theory with time: An integrated approach. *Academy of Management Review*, 29(3): 404-422.
- Gaimon, C. 1997. Planning information technology-knowledge worker systems. *Management Science*, 43(9): 1308-1328.
- Gaimon, C., and Bailey, J. 2012. Knowledge management for the entrepreneurial venture. *Production and Operations Management*. Forthcoming.
- Gavetti, G., and Levinthal, D. 2000. Looking forward and looking backward: Cognitive and experiential search. *Administrative Science Quarterly*, 45(1): 113-137.
- Gilson, L. L., Mathieu, J. E., Shalley, C. E., and Ruddy, T. M. 2005. Creativity and standardization: complementary or conflicting drivers of team effectiveness? *Academy of Management Journal*, 48(3): 521-531.
- Girotra, K., Terwiesch, C., and Ulrich, K. T. 2007. Valuing R&D projects in a portfolio: Evidence from the pharmaceutical industry. *Management Science*, 53(9): 1452-1466.
- Girotra, K., Terwiesch, C., and Ulrich, K. T. 2010. Idea generation and the quality of the best idea. *Management Science*, 56(4): 591-605.
- Gnyawali, D. R., and Park, B.-J. R. 2011. Co-opetition between giants: Collaboration with competitors for technological innovation. *Research Policy*, 40(5): 650-663.
- Griffin, A., and Page, A. L. 1993. An interim report on measuring product development success and failure. *Journal of Product Innovation Management*, 10(4): 291-308.
- Gruber, M., Harhoff, D., and Hoisl, K. 2012. Knowledge recombination across technological boundaries: Scientists vs. engineers. *Management Science*. Articles in Advance.
- Gupta, A. K., Smith, K. G., and Shalley, C. E. 2006. The interplay between exploration and exploitation. *The Academy of Management Journal*, 49(4): 693-706.
- Hall, B. H., Jaffe, A. B., and Trajtenberg, M. 2001. The NBER patent citation data file: Lessons, insights and methodological tools: *National Bureau of Economic Research*.
- Hamel, G., Doz, Y. L., and Prahalad, C. K. 1989. Collaborate with your competitors and win. *Harvard Business Review*, 67(1), 133-139.

- Hamel, G. 1991. Competition for competence and interpartner learning within international strategic alliances. *Strategic Management Journal*, 12(S1): 83-103.
- Hamilton, W. F., and Singh, H. 1992. The evolution of corporate capabilities in emerging technologies. *Interfaces*, 22(4): 13-23.
- Hartl, R., and Sethi, S. 1984. Optimal control of a class of systems with continuous lags: Dynamic programming approach and economic interpretations. *Journal of optimization Theory and Applications*, 43(1): 73-88.
- Haunschild, P. R., and Sullivan, B. N. 2002. Learning from complexity: Effects of prior accidents and incidents on airlines' learning. *Administrative Science Quarterly*, 47(4): 609-643.
- He, Z. L., and Wong, P. K. 2004. Exploration vs. exploitation: An empirical test of the ambidexterity hypothesis. *Organization Science*, 15(4): 481-494.
- Hillson, D. 2002. Extending the risk process to manage opportunities. *International Journal of Project Management*, 20(3): 235-240.
- Hoang, H., and Rothaermel, F. T. 2010. Leveraging internal and external experience: exploration, exploitation, and R&D project performance. *Strategic Management Journal*, 31(7): 734-758.
- Hora, M., and Dutta, D. K. 2012. Entrepreneurial firms and downstream alliance partnerships: Impact of portfolio depth and scope on technology innovation and commercialization success. *Production and Operations Management*. Forthcoming.
- Huchzermeier, A., Loch, C. H. 2001. Project management under risk: Using the real options approach to evaluate flexibility in R&D. *Management Science*, 47(1): 85.
- Im, G., and Rai, A. 2008. Knowledge sharing ambidexterity in long-term interorganizational relationships. *Management Science*, 54(7): 1281-1296.
- Inkpen, A. C. 2000. Learning through joint ventures: a framework of knowledge acquisition. *Journal of Management Studies*, 37(7): 1019-1044.
- Ittner, C. D., Nagar, V., and Rajan, M. V. 2001. An empirical examination of dynamic quality-based learning models. *Management Science*, 47(4): 563-578.
- Jain, A. 2013. Learning by doing and the locus of innovative capability in biotechnology research. *Organization Science*. Forthcoming.
- Jain, A., and Kogut, B. 2013. Memory and organizational evolvability in a neutral landscape. *Organization Science*. Forthcoming.

- Jorgensen, S. 1982. A survey of some differential games in advertising. *Journal of Economic Dynamics and Control*, 4: 341-369.
- Kamien, M. I., and Schwartz, N. L. 1975. Market structure and innovation: A survey. *Journal of Economic Literature*, 13(1): 1-37.
- Kane, V. E. 1986. Process capability indices. *Journal of Quality Technology*, 18(1): 41-52.
- Kanter, R. M. 1994. Collaborative advantage: the art of alliances. *Harvard Business Review*, 72(4): 96-108.
- Katila, R., and Ahuja, G. 2002. Something old, something new: A longitudinal study of search behavior and new product introduction. *Academy of Management Journal*: 45(6):1183-1194.
- Kavadias, S., and Sommer, S. C. 2009. The effects of problem structure and team diversity on brainstorming effectiveness. *Management Science*, 55(12): 1899-1913.
- KC, D., Staats, B. R., and Gino, F. 2013. Learning from my success and from others' failure: Evidence from minimally invasive cardiac surgery. *Management Science*. Forthcoming.
- Kelm, K. M., Narayanan, V., and Pinches, G. 1995. Shareholder value creation during R&D innovation and commercialization stages. *Academy of Management Journal*, 38(3): 770-786.
- Kerber, R. 2004. The making of a blockbuster. How Boston Scientific gambled on its new Taxus stent. *The Boston Globe*. April 26, 2004. Retrieved from [http://www.boston.com/business/technology/biotechnology/articles/2004/04/26/the\\_making\\_of\\_a\\_blockbuster](http://www.boston.com/business/technology/biotechnology/articles/2004/04/26/the_making_of_a_blockbuster).
- Khanna, T. 1995. Racing behavior technological evolution in the high-end computer industry. *Research Policy*, 24(6): 933-958.
- Khanna, T., Gulati, R., and Nohria, N. 1998. The dynamics of learning alliances: competition, cooperation, and relative scope. *Strategic Management Journal*, 19(3): 193-210.
- Kim, J. Y., Kim, J. Y. J., and Miner, A. S. 2009. Organizational learning from extreme performance experience: The impact of success and recovery experience. *Organization Science*, 20(6): 958-978.
- Klevorick, A. K., Levin, R. C., Nelson, R. R., and Winter, S. G. 1995. On the sources and significance of interindustry differences in technological opportunities. *Research Policy*, 24(2): 185-205.



- Koput, K. W. 1997. A chaotic model of innovative search: some answers, many questions. *Organization Science*, 8(5):528-542.
- Koza, M., and Lewin, A. 2000. Managing partnerships and strategic alliances: raising the odds of success. *European Management Journal*, 18(2): 146-151.
- Krishnan, V., Eppinger, S. D., and Whitney, D. E. 1997. A model-based framework to overlap product development activities. *Management Science*, 43(4): 437-451.
- Krishnan, V., and Loch, C. H. 2005. A retrospective look at production and operations management articles on new product development. *Production and Operations Management*, 14(4): 433-441.
- Lapr , M. A., and Nembhard, I. M. 2010. Inside the organizational learning curve: Understanding the organizational learning process. *Foundations and Trends in Technology, Information and Operations Management*, 4(1): 1-103.
- Lavie, D., and Rosenkopf, L. 2006. Balancing exploration and exploitation in alliance formation. *Academy of Management Journal*, 49(4): 797.
- Lavie, D., Stettner, U., and Tushman, M. L. 2010. Exploration and exploitation within and across organizations. *The Academy of Management Annals*, 4(1): 109-155.
- Lazer, D., and Friedman, A. 2007. The network structure of exploration and exploitation. *Administrative Science Quarterly*, 52(4): 667-694.
- Lee, D., and Van den Steen, E. 2010. Managing know-how. *Management Science*, 56(2): 270-285.
- Leiponen, A., and Helfat, C. E. 2010. Innovation objectives, knowledge sources, and the benefits of breadth. *Strategic Management Journal*, 31(2): 224-236.
- Lenfle, S., and Loch, C. 2010. Lost roots: how project management came to emphasize control over flexibility and novelty. *California Management Review*, 53(1): 32-55.
- Leonard-Barton, D. 1992. Core capabilities and core rigidities: A paradox in managing new product development. *Strategic Management Journal*, 13(S1): 111-125.
- Leonardi, P. M. 2011. Early prototypes can hurt a team's creativity. *Harvard Business Review*. 89(12):28.
- Lerner, J. 1997. An empirical exploration of a technology race. *The Rand Journal of Economics* 28(2): 228-247.

- Levinthal, D. A. 1997. Adaptation on rugged landscapes. *Management Science*, 43(7): 934-950.
- Levinthal, D., and March, J. 1993. The myopia of learning. *Strategic Management Journal*, 14(S2): 95-112
- Levitt, B., and March, J. G. 1988. Organizational learning. *Annual Review of Sociology*. 14 319-340.
- Libby, R., and Fishburn, P. C. 1977. Behavioral models of risk taking in business decisions: A survey and evaluation. *Journal of Accounting Research*, 15(2): 272-292.
- Linton, J. D., and Walsh, S. T. 2004. Integrating innovation and learning curve theory: an enabler for moving nanotechnologies and other emerging process technologies into production. *R&D Management*, 34(5): 517-526.
- Loch, C. H., and Terwiesch, C. 1998. Communication and uncertainty in concurrent engineering. *Management Science*, 44(8): 1032-1048.
- Loch, C. H., Terwiesch, C., and Thomke, S. 2001. Parallel and sequential testing of design alternatives. *Management Science*, 47(5): 663-678.
- Loch, C. H., and Tapper, U. 2002. Implementing a strategy-driven performance measurement system for an applied research group. *Journal of Product Innovation Management*, 19(3): 185-198.
- Loebecke, C., Van Fenema, P. C., and Powell, P. 1999. Co-opetition and knowledge transfer. *ACM SIGMIS Database*, 30(2): 14-25.
- Luo, L., Kannan, P., Besharati, B., and Azarm, S. 2005. Design of robust new products under variability: marketing meets design. *Journal of Product Innovation Management*, 22(2): 177-192.
- Madsen, P. M., and Desai, V. 2010. Failing to learn? The effects of failure and success on organizational learning in the global orbital launch vehicle industry. *The Academy of Management Journal (AMJ)*, 53(3): 451-476.
- Manso, G. 2011. Motivating innovation. *The Journal of Finance*, 66(5): 1823-1860.
- March, J. G. 1991. Exploration and exploitation in organizational learning. *Organization Science* 2(1): 71-87.
- March, J. G. 2003. Understanding organizational adaptation. *Society and Economy*, 25(1): 1-10.

- March, J. G., and Shapira, Z. 1992. Variable risk preferences and the focus of attention. *Psychological Review* 99(1): 172.
- McGill, J. P., and Santoro, M. D. 2009. Alliance portfolios and patent output: The case of biotechnology alliances. *Engineering Management, IEEE Transactions* 56(3): 388-401.
- McGrath, R. G. 1999. Falling forward: Real options reasoning and entrepreneurial failure. *Academy of Management Review*, 24(1): 13-30.
- McKee, D. 1992. An organizational learning approach to product innovation. *Journal of Product Innovation Management*, 9(3): 232-245.
- Müller, D. 2010. Alliance coordination, dysfunctions, and the protection of idiosyncratic knowledge in strategic learning alliances. Working Paper. Leibniz Information Centre for Economics.
- Narayanan, S., Balasubramanian, S., and Swaminathan, J. M. 2009. A matter of balance: Specialization, task variety, and individual learning in a software maintenance environment. *Management Science*, 55(11): 1861-1876.
- Nelson, R. R. 2008. Bounded rationality, cognitive maps, and trial and error learning. *Journal of Economic Behavior & Organization*, 67(1): 78-89.
- Nerkar, A. 2003. Old is gold? The value of temporal exploration in the creation of new knowledge. *Management Science*, 49(2): 211-229.
- Obstfeld, D. 2012. Creative projects: a less routine approach toward getting new things done. *Organization Science*, 23(6): 1571-1592.
- Oraiopoulos, N., and Kavadias, S. 2008. The role of informational spillovers on competitive R&D Search. Georgia Institute of Technology Working Paper.
- O'Reilly III, C. A., and Tushman, M. L. 2008. Ambidexterity as a dynamic capability: Resolving the innovator's dilemma. *Research in Organizational Behavior*, 28: 185-206.
- O'Reilly, C., and Tushman, M. 2013. Organizational ambidexterity: Past, present and future. *The Academy of Management Perspectives*. amp. 2013.0025
- Oxley, J. E., and Sampson, R. C. 2004. The scope and governance of international R&D alliances. *Strategic Management Journal*, 25(8-9): 723-749.
- Ozkan, G., Gaimon, C., and Kavadias, S. 2009. Knowledge management strategies for product and process design teams. Georgia Institute of Technology Working Paper.

- Petkova, A. P. 2009. A theory of entrepreneurial learning from performance errors. *International Entrepreneurship and Management Journal*, 5(4): 345-367.
- Pisano, G. P. 1994. Knowledge, integration, and the locus of learning: An empirical analysis of process development. *Strategic Management Journal*, 15(S1): 85-100.
- Quintana-García, C., and Benavides-Velasco, C. A. 2008. Innovative competence, exploration and exploitation: The influence of technological diversification. *Research Policy*, 37(3): 492-507.
- Rahmandad, H. 2008. Effect of delays on complexity of organizational learning. *Management Science*, 54(7): 1297-1312.
- Rajgopal, S., and Shevlin, T. 2002. Empirical evidence on the relation between stock option compensation and risk taking. *Journal of Accounting and Economics*, 33(2): 145-171.
- Raisch, S., Birkinshaw, J. 2008. Organizational ambidexterity: Antecedents, outcomes, and moderators. *Journal of Management*, 34(3): 375.
- Raisch, S., Birkinshaw, J., Probst, G., and Tushman, M. L. 2009. Organizational ambidexterity: Balancing exploitation and exploration for sustained performance. *Organization Science*, 20(4): 685-695.
- Reber, A. S. 1989. Implicit learning and tacit knowledge. *Journal of Experimental Psychology*, 118(3): 219.
- Reinganum, J. F. 1981. Dynamic games of innovation. *Journal of Economic Theory*, 25(1): 21-41.
- Rerup, C., and Feldman, M. S. 2011. Routines as a source of change in organizational schemata: The role of trial-and-error learning. *Academy of Management Journal*, 54(3): 577-610.
- Roels, G., and Su, X. 2013. Optimal Design of Social Comparison Effects: Setting Reference Groups and Reference Points. *Management Science*. Forthcoming.
- Rosenkopf, L., and Nerkar, A. 2001. Beyond local search: boundary-spanning, exploration, and impact in the optical disk industry. *Strategic Management Journal*, 22(4): 287-306.
- Rothaermel, F. T., and Alexandre, M. T. 2009. Ambidexterity in technology sourcing: The moderating role of absorptive capacity. *Organization Science*, 20(4): 759-780.

- Rothaermel, F. T., and Deeds, D. L. 2004. Exploration and exploitation alliances in biotechnology: A system of new product development. *Strategic Management Journal*, 25(3): 201-221.
- Schilling, M. A., and Green, E. 2011. Recombinant search and breakthrough idea generation: An analysis of high impact papers in the social sciences. *Research Policy*, 40(10): 1321-1331.
- Schilling, M. A., Vidal, P., Ployhart, R. E., and Marangoni, A. 2003. Learning by doing something else: Variation, relatedness, and the learning curve. *Management Science*, 49(1): 39-56.
- Schulz, M. 2001. The uncertain relevance of newness: Organizational learning and knowledge flows. *Academy of Management Journal*, 44(4): 661-681.
- Sethi, R., and Iqbal, Z. 2008. Stage-gate controls, learning failure, and adverse effect on novel new products. *Journal of Marketing*, 72(1): 118-134.
- Sethi S. P., G. L. Thompson. 2000. Optimal control theory: Applications to management science and economics. Kluwer Academic Publishers, Boston, MA.
- Shan, W. 1990. An empirical analysis of organizational strategies by entrepreneurial high-technology firms. *Strategic Management Journal*, 11(2): 129-139.
- Shane, S. 2000. Prior knowledge and the discovery of entrepreneurial opportunities. *Organization Science*, 11(4): 448-469.
- Siggelkow, N., and Levinthal, D. A. 2003. Temporarily divide to conquer: Centralized, decentralized, and reintegrated organizational approaches to exploration and adaptation. *Organization Science*, 14(6): 650-669.
- Singh, J., and Fleming, L. 2010. Lone inventors as sources of breakthroughs: Myth or reality? *Management Science*, 56(1): 41-56.
- Sitkin, S. B. 1992. Learning through failure: The strategy of small losses. *Research in Organizational Behavior*, 14: 231-231.
- Sitkin, S. B., and Pablo, A. L. 1992. Reconceptualizing the determinants of risk behavior. *Academy of Management Review*, 17(1): 9-38.
- Smith, W. K., Binns, A., and Tushman, M. L. 2010. Complex business models: Managing strategic paradoxes simultaneously. *Long Range Planning*, 43(2): 448-461.
- Sommer, S. C., and Loch, C. H. 2009. Incentive contracts in projects with unforeseeable uncertainty. *Production and Operations Management*, 18(2): 185-196.

- Sood, A., and Tellis, G. J. 2009. Do innovations really pay off? Total stock market returns to innovation. *Marketing Science*, 28(3): 442-456.
- Sørensen, J. B., and Stuart, T. E. 2000. Aging, obsolescence, and organizational innovation. *Administrative Science Quarterly*, 45(1): 81-112.
- Sorenson, O., Rivkin, J. W., and Fleming, L. 2006. Complexity, networks and knowledge flow. *Research Policy*, 35(7), 994-1017.
- Sorescu, A. B., Chandy, R. K., and Prabhu, J. C. 2003. Sources and financial consequences of radical innovation: Insights from pharmaceuticals. *Journal of Marketing*, 67: 82-102.
- Spencer, J. W. 2003. Firms' knowledge-sharing strategies in the global innovation system: empirical evidence from the flat panel display industry. *Strategic Management Journal*, 24(3): 217-233.
- Staw, B. M., Sandelands, L. E., Dutton, J. E. 1981. Threat rigidity effects in organizational behavior: A multilevel analysis. *Administrative Science Quarterly*, 26(4): 501-524.
- Taylor, A., and Greve, H. 2006. Superman or the fantastic four? Knowledge combination and experience in Innovative Teams. *Academy of Management Journal*, 49(4): 723-740.
- Teece, D. J. 1992. Competition, cooperation, and innovation:: Organizational arrangements for regimes of rapid technological progress. *Journal of Economic Behavior and Organization*, 18(1): 1-25.
- Terwiesch, C., and Ulrich, K. T. 2009. Innovation tournaments: Creating and selecting exceptional opportunities. Harvard Business School Press.
- Thomke, S. H. 1998. Managing experimentation in the design of new products. *Management Science*: 44(6): 743-762.
- Thomke, S. 2001. Enlightened experimentation: The new imperative for innovation. *Harvard Business Review*, 79(2): 66-75.
- Thomke, S., and Bell, D. E. 2001. Sequential testing in product development. *Management Science*, 47(2): 308-323.
- Thomke, S., Von Hippel, E., and Franke, R. 1998. Modes of experimentation: an innovation process and competitive variable. *Research Policy*, 27(3): 315-332.

- Thomke, S., and Fujimoto, T. 2000. The effect of "front-loading" problem-solving on product development performance. *Journal of Product Innovation Management*, 17(2): 128-142.
- Thornhill, S., and White, R. E. 2007. Strategic purity: A multi-industry evaluation of pure vs. hybrid business strategies. *Strategic Management Journal*, 28(5): 553-561.
- Tian, X., and Wang, T. 2011. Tolerance for failure and corporate innovation. *Review of Financial Studies*. Forthcoming.
- Trajtenberg, M. 1990. A penny for your quotes: patent citations and the value of innovations. *The Rand Journal of Economics*, 21(1): 172-187.
- Tsetlin, I., Gaba, A., and Winkler, R. L. 2004. Strategic choice of variability in multi-round contests and contests with handicaps. *Journal of Risk and Uncertainty*, 29(2): 143-158.
- Tucker, A. L., Nemhard, I. M., and Edmondson, A. C. 2007. Implementing new practices: An empirical study of organizational learning in hospital intensive care units. *Management Science*, 53(6): 894-907.
- Tushman, M. L., and Anderson, P. 1986. Technological discontinuities and organizational environments. *Administrative Science Quarterly*, 31(3):439-465.
- Tushman, M. L., and O'Reilly III, C. A. 2006. Ambidextrous organizations: Managing evolutionary and revolutionary change. *California Management Review*, 38(4):8-30.
- Tyebjee, T. T., and Bruno, A. V. 1984. A model of venture capitalist investment activity. *Management Science*, 30(9): 1051-1066.
- Van De Ven, A. H., and Polley, D. 1992. Learning while innovating. *Organization Science*, 7(6): 92-116.
- Vermeulen, F., and Barkema, H. 2001. Learning through acquisitions. *Academy of Management Journal*, 44(3), 457-476.
- Von Hippel, E., and Von Krogh, G. 2003. Open source software and the "private-collective" innovation model: Issues for organization science. *Organization Science*, 14(2): 209-223.
- Von Hippel, E., and Von Krogh, G. 2006. Free revealing and the private-collective model for innovation incentives. *R&D Management*, 36(3): 295-306.
- Walley, K. 2007. Coopetition: An introduction to the subject and an agenda for research. *International Studies of Management and Organization*, 37(2): 11-31.

- Ward, T. B. 2004. Cognition, creativity, and entrepreneurship. *Journal of Business Venturing*, 19(2): 173-188.
- Wiseman, R. M., Bromiley, P. 1996. Toward a model of risk in declining organizations: An empirical examination of risk, performance and decline. *Organization Science*, 7(5): 524.
- West, J., and Iansiti, M. 2003. Experience, experimentation, and the accumulation of knowledge: the evolution of R&D in the semiconductor industry. *Research Policy*, 32(5): 809-825.
- Wong, S.S. 2004. Distal and local group learning: Performance trade-offs and tensions. *Organization Science*, 15(6): 645-656.
- Wu, B., and Knott, A. M. 2006. Entrepreneurial risk and market entry. *Management Science*, 52(9):1315-1330.
- Wu, J., and Shanley, M. T. 2009. Knowledge stock, exploration, and innovation: Research on the United States electromedical device industry. *Journal of Business Research*, 62(4): 474-483.
- Xiao, W., and Gaimon, C. 2012. Knowledge Creation and Knowledge Transfer in New Product Development Projects. Georgia Institute of Technology Working Paper.
- Xue, Y. 2007. Make or buy new technology: The role of CEO compensation contract in a firm's route to innovation. *Review of Accounting Studies*, 12(4): 659-690.
- Yang, H., Phelps, C., and Steensma, H. K. 2010. Learning from what others have learned from you: The effects of knowledge spillovers on originating firms. *Academy of Management Journal*, 53(2), 371-389.
- Yelle, L. E. 1979. The learning curve: Historical review and comprehensive survey. *Decision Sciences*, 10(2): 302-328.
- Zellner, A. 1962. An efficient method of estimating seemingly unrelated regressions and tests for aggregation bias. *Journal of the American Statistical Association*, 57(298): 348-368.
- Zhou, Y. M. 2011. Synergy, coordination costs, and diversification choices. *Strategic Management Journal*, 32(6): 624-639.