



Entrepreneurial Strategy and Learning

Citation

Peterson, Aticus Arthur. 2025. Entrepreneurial Strategy and Learning. Doctoral Dissertation, Harvard University Graduate School of Arts and Sciences.

Link

<https://dash.harvard.edu/handle/1/42719621>

Terms of use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material (LAA), as set forth at

<https://harvardwiki.atlassian.net/wiki/external/NGY5NDE4ZjgzNTc5NDQzMGIzZWZhMGFIOWI2M2EwYTg>

Accessibility

<https://accessibility.huit.harvard.edu/digital-accessibility-policy>

Share Your Story

The Harvard community has made this article openly available. Please share how this access benefits you. [Submit a story](#)

HARVARD UNIVERSITY
Graduate School of Arts and Sciences



DISSERTATION ACCEPTANCE CERTIFICATE

The undersigned, appointed by the Committee for the PhD in Business Administration have examined a dissertation entitled

Entrepreneurial Strategy and Learning

Presented by **Aticus Peterson**

candidate for the degree of Doctor of Philosophy and hereby certify that it is worthy of acceptance.

Signature

A handwritten signature in black ink, appearing to read 'Andy Wu', written over a horizontal line.

Andy Wu, Chair

Signature

A handwritten signature in black ink, appearing to read 'Dennis Yao', written over a horizontal line.

Dennis Yao

Signature

A handwritten signature in black ink, appearing to read 'Eric Van den Steen', written over a horizontal line.

Eric Van den Steen

Signature

A handwritten signature in black ink, appearing to read 'Jan Rivkin', written over a horizontal line.

Jan Rivkin

Signature

A handwritten signature in black ink, appearing to read 'Rory McDonald', written over a horizontal line.

Rory McDonald

Date: April 4, 2025

Entrepreneurial Strategy and Learning

A DISSERTATION PRESENTED
BY
ATICUS PETERSON
TO
THE DEPARTMENT OF STRATEGY

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY
IN THE SUBJECT OF
BUSINESS ADMINISTRATION

HARVARD UNIVERSITY
CAMBRIDGE, MASSACHUSETTS
APRIL 2025

©2025 – ATICUS PETERSON
ALL RIGHTS RESERVED.

Entrepreneurial Strategy and Learning

ABSTRACT

Learning is a cornerstone of competitive advantage, especially in entrepreneurial contexts. Despite extensive research on entrepreneurial learning methodologies, we still lack understanding of fundamental constraints that limit learning effectiveness. This dissertation investigates key structural obstacles that hinder effective entrepreneurial learning: resource constraints, uncertainty, interdependencies, and bounded rationality. It provides novel insights into how these barriers manifest and interact in entrepreneurial contexts. The first study reveals a counterintuitive relationship between experience and strategic foresight: as entrepreneurs add features to successive products, the total interdependencies grow faster than their ability to anticipate them, leading to increasingly inaccurate forecasts. The second study reveals how novel recombination, while often associated with opportunities for outsized returns, also incurs the cost of managing novel complexity—obscuring the time and resources required for execution and increasing the risk of bridge financing and shutdown. The third study identifies a “mediocrity trap,” in which firms launching lower-quality products receive ambiguous feedback and, as a result, persist longer despite negative market signals. This study challenges assumptions about learning with minimum viable products and rapid pivoting. Together, these studies highlight critical tensions in entrepreneurial learning and demonstrate that many learning challenges stem from structural barriers rather than purely cognitive limitations. This integrated perspective contributes to scholarly understanding of entrepreneurial strategy while offering practical approaches for entrepreneurs seeking to navigate these barriers and learn under constraint.

Contents

Title Page	i
Copyright	ii
Abstract	iv
TABLE OF CONTENTS	iv
LIST OF FIGURES	vi
LIST OF TABLES	viii
AUTHORLIST	ix
ACKNOWLEDGMENTS	x
INTRODUCTION	i
1 ENTREPRENEURIAL LEARNING AND STRATEGIC FORESIGHT	8
1.1 Introduction	8
1.2 Theoretical Background	12
1.3 Hypothesis Development	18
1.4 Empirical Methods	24
1.5 Results	35
1.6 Discussion and Conclusion	43
2 COMPLEXITY AND TIMING IN NOVEL VENTURES	50
2.1 Introduction	50
2.2 Theoretical Background	54
2.3 Hypothesis Development	62
2.4 Empirical Methods	68
2.5 Results	79
2.6 Discussion	87
2.7 Conclusion	94

3	THE MEDIOCRITY TRAP IN ENTREPRENEURIAL LEARNING	95
3.1	Introduction	95
3.2	Theoretical Background	98
3.3	Simple Model	103
3.4	Hypothesis Development	112
3.5	Empirical Methods	116
3.6	Results	126
3.7	Discussion	130
	APPENDIX A ENTREPRENEURIAL LEARNING AND STRATEGIC FORESIGHT	134
	A.1 Online Appendix	134
	APPENDIX B COMPLEXITY AND TIMING IN NOVEL VENTURES	182
	B.1 Appendix	182
	APPENDIX C THE MEDIOCRITY TRAP IN ENTREPRENEURIAL LEARNING	197
	C.1 Appendix	197
	BIBLIOGRAPHY	249

Listing of figures

1.1	Conceptual Model of Experience and Interdependencies.	16
1.2	Example Products by Entrepreneur Over Time.	25
1.3	Density Plot of <i>Actual Time</i> and <i>Predicted Time</i>	34
1.4	<i>Actual Time</i> and <i>Predicted Time</i> with Experience.	39
2.1	Novel Recombination and the Timing of Value Realization.	65
2.2	Novel Recombination and Valuation Over Time.	82
2.3	Novel Recombination and Time to Catch-up.	88
3.1	Model Overview.	104
3.2	Updated Prior Belief μ_1 After Failure to Sell.	107
3.3	Model: Illustrative Case.	110
A.1	Plum Geek Robotics Product Features.	176
A.2	Project Complexity Examples.	177
A.3	Variable Distributions.	178
A.4	Outcome–Aspiration Gap and Accuracy of Strategic Foresight.	179
A.5	Deviation and Delay.	180
A.6	Deviation and Unforeseen Interdependencies.	181
C.1	Binned Scatter of Salespeople and Paid Keywords.	222

List of Tables

1.1	Summary Statistics.	32
1.2	Pairwise Correlation Matrix of Independent Variables.	33
1.3	Features and Unforeseen Interdependencies.	36
1.4	Delivery and Delay.	38
2.1	Summary Statistics - Firm-Year Level.	76
2.2	Summary Statistics - Firm-Round Level.	78
2.3	Novel Recombination and Complexity.	80
2.4	Novel Recombination and Valuation Over Time.	83
2.5	Novel Recombination and Time to High Valuation.	84
2.6	Novel Recombination and Performance.	85
3.1	Summary Statistics.	122
3.2	Correlation Table.	123
3.3	Tendency to Persist and Sell Again.	127
3.4	Tendency to Pivot or Shutdown.	128
3.5	Initial Sales and Product Investment and Long-Term Outcomes.	130
A.1	Qualitative Interview Sample.	136
A.2	Qualitative Interview Sample Summary Statistics.	136
A.3	Project Reviewer Backgrounds.	144
A.4	Variable Quantiles.	147
A.5	Unforeseen Interdependencies Alternative.	148
A.6	Non-Linear Analysis of Complexity and Unforeseen Interdependencies.	151
A.7	Non-Linear Analysis of Delivery and Delay.	152
A.8	Funding Deviation and Performance.	159
A.9	Predicted Time and Peer Group Comparison.	162
A.10	Funding and Peer Group Comparison.	164
A.11	Consequences of Delay.	167
A.12	Features and Negative Comments.	168
A.13	Creator Engagement.	170

A.14	Incentive to Overpromise.	172
A.15	Exiting After VC Financing.	173
A.16	Exiting After Delay.	175
B.1	Clustering Method Performance Summary.	185
B.2	Summary Statistics by Clustering Method.	186
B.3	Pairwise Correlation Matrix by Clustering Method.	186
B.4	Novel Recombination and Entry Order.	188
B.5	Tech Stack Summary Statistics and Pairwise Correlation.	190
B.6	Novel Recombination and Tech Stack Innovation.	190
B.7	Firm-Year Panel Pairwise Correlation Matrix.	193
B.8	Firm-Round Panel Pairwise Correlation Matrix.	194
B.9	Software Firm-Month Panel Summary Statistics and Pairwise Correlation.	195
B.10	Novel Recombination and Modularity.	196
C.1	Alternative Measures of Persistence in Selling Again.	209
C.2	BuiltWith Selection: Post-2012.	211
C.3	Persistence in Selling Again and Competition Variables.	216
C.4	Persistence in Selling Again and Competition Fixed Effects.	217
C.5	Single vs Multiple Products Over Time.	220
C.6	Paid Keywords.	224
C.7	Strategic Shifts Via CEO Change.	226
C.8	PageRank Analysis: Impact of Customer Size.	228
C.9	Leading Customer Acquisition Cost.	232
C.10	Sell Again: Two and Three Year Lead.	234

Chapter 1 and Chapter 2 are co-authored with Andy Wu. Chapter 3 is sole-authored.

Acknowledgments

I'm grateful for this chance to reflect on all those who made this dissertation possible. Some contributed directly with ideas, feedback, and analysis; others provided encouragement and support throughout the process; and many more shaped who I am today in ways that made this whole journey possible. This dissertation reflects the goodness and brilliance of the many people who have built and shaped me throughout my life.

First, I want to express my profound gratitude to my dissertation committee:

Andy Wu, my chair, who I've worked with from the very beginning: thank you for the mentorship, for freely giving of your time and ideas, for modeling how to collaborate, for always expressing appreciation for work done, for listening and valuing my ideas.

Dennis Yao: thank you for all your thoughtful comments, for carefully annotating my papers with your input, for inspiring me with your unparalleled command of the literature and ability to make connections across ideas.

Eric Van den Steen: thank you for opening the world of modeling to me, for being so incredibly thoughtful and encouraging, for your patience as I learn, for helping me build a foundation of skills and understand the principles.

Jan Rivkin: thank you for all the walks across our (not so rugged) campus landscape, for asking precise questions to get to the core and move things forward, for helping and inspiring me to not just do what is right but what is most right.

Rory McDonald: thank you for all the encouragement, for guiding and teaching me to place my work in the strategy and entrepreneurship literatures, for the guidance on qualitative work, for the personal support.

I am grateful to the broader Strategy Unit for the community and support. Specifically to Juan Alcacer (for teaching me the foundations), to Hong Luo (for guiding us in RDC), to Rem Koning (for the advice and encouragement), and to Becky Karp (for the real conversations and many projects). I also am extremely grateful to Roy Wang and Bailey Hughes for helping me schedule so many meetings and all the email back and forth.

My cohort made the journey fun. Thank you to my horizontal cohort, including Botir Kobilov and Natalie Epstein (together, the Party Planning Committee), Laura Katsnelson, Daniel Yue, Caleb Kwon, Maya Balakrishnan, Kunal Sangani, and Jimin Nam. And to my vertical cohort, including Nataliya Wright, Innessa Colaiacova, Dan Brown, Chris Eaglin, Liz Calder, and Haiyang Zhang. And to Michael Christensen. To all of you, your feedback and support were incredibly helpful, but your

friendships are invaluable.

My thanks to Ronald Wang for the incredible, above-and-beyond research assistance. To HBS IT and the Baker Library team. To HBS Dining for so many meals. To the Faculty Commons for fueling me with Diet Coke. And to Dinorah for always saying hi and stopping to chat.

The Doctoral Programs Office provided unparalleled guidance, support, and encouragement. Thank you, Jen—I'll always remember you showing up to the hospital with gift baskets in hand. I'm also grateful to Angela, Marais, Joel, Finley, Darlene, and Aric.

Personally, I am grateful for incredible friends in the LDSSA over the years. I loved the exponential conversations on walk-and-talks with Kirk Allen. I felt so supported and encouraged by my wonderful neighbors at 993, Judith and Josie. And thank you to Kevin Keovongsa for helping me through hard times.

This feels like the culmination of so many years of education and the inspiring, selfless teachers along the way. Thank you, Beth Gentile, Grace Martin, Kim Allen, Jen Flitton, Vicki Bowling, Zara Westbrook, Julie Nelson, and Margo Levine.

And of course thank you to my foundation, my family. I am so blessed to have such loving and supporting parents, Cynthia and Bruce, who have inspired and encouraged me from day one. Anna and Amadeus are the best siblings one could want, and I'm excited to welcome Taylor. I'm grateful to my grandparents, especially my Grandpa Gary Sandquist, who instilled the power of learning in our whole family and helped support my education. Thank you to my in-laws: Melinda, Collin, Allie, and Duncan. Thank you to my dog, Nigel, who joined me the month before I started my PhD journey and always makes me smile.

Victoria: I love you and am so grateful for all your love and support. You inspire me. Griffin (and any future children): I love you and, in many ways, do this for you. I hope you, like me, will be surrounded by people who love and support you. I'll always be one of them.

THIS DISSERTATION IS DEDICATED TO MY FAMILY.

Introduction

The entrepreneurial process fundamentally involves discovering and exploiting opportunities under conditions of uncertainty. This process requires continuous learning as entrepreneurs test hypotheses and refine their understanding of products, markets, and business models. As such, superior learning is recognized as a potential source of competitive advantage.

However, entrepreneurs face significant structural barriers to learning. While research has examined learning processes, less attention has been paid to the fundamental obstacles that impede this learning. Understanding these barriers is crucial for both theoretical frameworks that explain entrepreneurial strategy and practical approaches that improve outcomes.

This dissertation examines significant barriers to entrepreneurial learning, including uncertainty, interdependencies, bounded rationality, and resource constraints. Through three chapters addressing different aspects of these barriers, this work contributes to a more nuanced understanding of entrepreneurial strategy and learning.

BARRIERS TO ENTREPRENEURIAL LEARNING

RESOURCE CONSTRAINTS

Resource constraints present a pervasive barrier to learning in entrepreneurial contexts. These constraints fundamentally shape entrepreneurs' ability to learn effectively. For example, financial constraints limit the scope and scale of experiments entrepreneurs can conduct. Human resource constraints restrict the specialized expertise available to interpret information. Technological constraints limit access to tools that could enhance learning capabilities. These constraints are particularly acute for new ventures that lack the accumulated resources of established firms.

The scarcity of resources creates challenging tradeoffs that affect learning quality. Entrepreneurs must balance exploration (seeking new knowledge) and exploitation (refining existing knowledge), often under pressure to demonstrate short-term progress. Under severe limitations, entrepreneurs may resort to smaller sample sizes or simplified analytical techniques—leading to less reliable conclusions. Resource constraints interact with and often amplify the other barriers to learning.

This creates a fundamental tension: entrepreneurs need to learn efficiently to conserve scarce resources, yet these same resource constraints often make efficient learning more difficult. This tension shapes entrepreneurial strategy and outcomes in profound ways.

UNCERTAINTY

Uncertainty represents a fundamental barrier to entrepreneurial learning. Unlike risk, which involves known probabilities, true uncertainty involves unknown and unknowable elements—a defining characteristic of entrepreneurial environments.

This pervasive uncertainty complicates learning in several ways. It creates challenges for decision-making when entrepreneurs cannot assign probabilities to potential outcomes or even identify all

possible outcomes. It makes learning through experimentation difficult because entrepreneurs often cannot determine what knowledge will be valuable until after it's acquired. Entrepreneurs must learn without the benefit of historical precedents, established heuristics, or reliable forecasting methods.

Furthermore, uncertainty evolves throughout the entrepreneurial journey. As initial uncertainties are resolved through learning, new ones often emerge. This dynamic uncertainty landscape requires entrepreneurs to continuously adapt their learning approaches, often without recognizing how the learning environment has changed. The most skilled entrepreneurs develop strategies for reducing uncertainty systematically, but the fundamental challenge of learning under uncertainty remains a persistent barrier.

INTERDEPENDENCIES

Complexity—specifically the web of interdependencies between components, activities, or decisions—represents another significant barrier to entrepreneurial learning. Grounded in complex systems theory and further developed through NK models, interdependencies create non-linear relationships that make entrepreneurial learning particularly challenging.

Interdependencies manifest in multiple ways: technical interdependencies between product features, market interdependencies between customer segments and attributes, and strategic interdependencies between business model elements. These create non-linear relationships where small changes in one element can produce disproportionately large effects elsewhere—a manifestation of causal ambiguity.

Additionally, interdependencies often create rugged performance landscapes where local optimization leads to suboptimal global outcomes. The challenges can also intensify as ventures develop and scale, with interdependencies that were not apparent in early stages emerging later (e.g. unknown unknowns), further complicating learning.

BOUNDED RATIONALITY

Even if entrepreneurs could overcome uncertainty and complexity, they would still face cognitive limitations that constrain their ability to learn effectively. The concept of bounded rationality, central to the behavioral theory of the firm, highlights how human cognitive constraints limit perfect rationality and optimal decision-making.

Bounded rationality manifests in entrepreneurial learning through several mechanisms. Entrepreneurs face attention limitations—they cannot monitor all relevant variables simultaneously. They face memory limitations that affect their ability to retrieve relevant experiences. They also face processing limitations in their ability to analyze complex information and identify patterns.

These cognitive constraints lead to simplification strategies that can distort learning. Entrepreneurs rely on heuristics that, while generally useful, can produce systematic biases. They engage in satisficing behavior, settling for “good enough” solutions rather than optimal ones. They also tend to maintain cognitive consistency, potentially rejecting disconfirming information.

CHAPTER OVERVIEWS

This dissertation examines barriers to entrepreneurial learning through three complementary empirical studies. Each chapter addresses specific aspects of how entrepreneurs navigate uncertainty, manage interdependencies, overcome bounded rationality, and operate under resource constraints.

ENTREPRENEURIAL LEARNING AND STRATEGIC FORESIGHT

This chapter investigates a counterintuitive finding: entrepreneurs make less accurate predictions as they gain experience executing projects. Through a study of 314 entrepreneurs across 722 crowd-funded hardware technology projects, we find that entrepreneurs miss their predicted timeline by nearly six additional weeks on each successive project, despite allocating more time for later projects.

This chapter primarily addresses the interdependencies barrier by identifying how complexity increases: as entrepreneurs gain experience, they incorporate more features into subsequent products, introducing interdependencies that expand beyond what they can anticipate. The paper also highlights bounded rationality—despite direct experience with previous projects, entrepreneurs remain unable to fully anticipate cascading interdependencies when adding seemingly simple features. These interdependencies make accurate prediction increasingly difficult.

The study challenges the assumption that experience necessarily improves strategic foresight, suggesting instead that growing interdependencies can outpace learning capacity. It provides a nuanced explanation for entrepreneurial failure, highlighting how structural barriers to learning affect even experienced entrepreneurs.

COMPLEXITY AND TIMING IN NOVEL VENTURES

This chapter examines how novel recombination—combining existing knowledge in new ways—affects execution time and venture performance. Analyzing 31,450 firms across 94,549 funding rounds, we find that novel recombination significantly extends execution timelines and increases failure risk. Increasing novel recombination extends time to valuation milestones, increases the likelihood of requiring bridge financing, and raises shutdown probability.

This paper illustrates how novel recombination creates complex systems with numerous interdependencies that must be managed during execution. Resource constraints become particularly apparent in findings on bridge financing and failure, as novel ventures often require more time than anticipated to resolve interdependencies, exceeding initially allocated resources.

By highlighting relationships between novelty, complexity, timing, and venture outcomes, this paper provides a framework for understanding why promising ventures struggle to translate valuable ideas into functioning businesses within expected timeframes. It demonstrates how barriers to learning interact to create execution challenges beyond what entrepreneurs and investors typically antici-

pate.

THE MEDIOCRITY TRAP IN ENTREPRENEURIAL LEARNING

This chapter examines a critical decision entrepreneurs face: whether to try selling their existing product or invest more in product development first. Through modeling and analysis of 1,445 enterprise software companies, I find that firms going to market with lower-quality products persist longer despite negative market signals and ultimately fail more often. Contrary to expected benefits of fast market entry, these firms expend more resources and take more time overall.

This paper highlights how resource constraints force tradeoffs between product development and sales. Specifically, it illustrates how the interdependencies between product quality and customer feedback affects signal noise. With low product quality, negative market feedback becomes ambiguous—entrepreneurs cannot determine whether failure stems from lack of market demand or insufficient product development.

This paper challenges the presumed compatibility between developing minimum viable products and rapidly pivoting based on market feedback, perhaps most broadly popularized in the Lean Startup approach. Instead, minimal products may predispose the entrepreneur to persist rather than pivot upon failure.

CONCLUSION

These chapters highlight important tensions in entrepreneurial learning: between developing complex products versus simpler ones, between early market entry versus greater product development, and between exploring novel recombinations versus more familiar combinations. A key insight is that barriers to entrepreneurial learning are often structural rather than purely cognitive. While psychological factors play a role, many learning challenges stem from fundamental tensions in the en-

trepreneurial process that cannot simply be eliminated through better cognitive strategies.

This integrated perspective suggests that effective entrepreneurial learning requires explicitly addressing these structural barriers. Entrepreneurs need approaches that help them manage interdependencies, design experiments that reduce uncertainty, acknowledge bounded rationality, and allocate limited resources to maximize learning effectiveness. By identifying and analyzing these barriers, this dissertation contributes to both scholarly understanding of entrepreneurial strategy and practical approaches to entrepreneurial decision-making.

1

Entrepreneurial Learning and Strategic Foresight

1.1 INTRODUCTION

Strategic foresight is the ability to accurately predict the consequences of a strategy and in turn pursue a superior course of action to build competitive advantage (Ahuja et al., 2005, Gavetti & Menon,

2016). Heterogeneity in strategic foresight implicitly underlies many theories of competitive advantage (Csaszar & Laureiro-Martínez, 2018): by predicting the value that resources generate after being developed (Barney, 1986) and the attractiveness of potential opportunities (Porter, 1980), entrepreneurs with superior foresight can build competitive advantage. As such, understanding the antecedents of strategic foresight should shed critical light on the origins of competitive advantage (Csaszar, 2018), yet research on this topic has been sparse until recently. Recent studies find that individual cognition (Gary & Wood, 2011, Kapoor & Wilde, 2020) and organizational structure (Csaszar, 2012) are determinants of strategic foresight. In particular, recent work emphasizes the specific importance of strategic foresight for entrepreneurs as they formulate strategy (Eisenhardt & Bingham, 2017, Ott et al., 2017).

Although an entrepreneur may need strategic foresight to build competitive advantage, the conditions that enable an entrepreneur to build strategic foresight are unclear, particularly as she learns from past experience (Nelson & Winter, 2002) and applies it to the next entrepreneurial opportunity (Gavetti, 2012). An extensive body of work on serial entrepreneurship documents how entrepreneurs can improve performance from venture to venture (e.g., Gompers et al., 2010, Stuart & Abetti, 1990), but that improvement can arise from a variety of factors besides improved foresight. Through experience, a serial entrepreneur accumulates a multitude of advantages: resource access and relationships (Clough et al., 2018, Hsu, 2007), opportunities (Gruber et al., 2008), and knowledge about the consequences of past decisions (Minniti & Bygrave, 2001, Paik, 2014). Although most of this work does not specifically speak to whether experience improves the strategic foresight of the entrepreneur, it generally makes this implicit assumption. However, recent work calls this untested assumption into question altogether. Cognitive limitations constrain what an entrepreneur can learn from experience and use for effective judgement (Cassar, 2014, Cohen et al., 2019). In an important study, Eggers & Song (2015) demonstrate that boundedly rational entrepreneurs may misattribute the sources of past performance. Thus, it remains an open question whether entrepreneurs learn from experience

to improve strategic foresight.

This study explores this open question for entrepreneurial strategy. Specifically, how does experience from executing past projects affect the accuracy of an entrepreneur's strategic foresight on the subsequent project? Given that tackling complexity is a defining characteristic of the phenomenon of entrepreneurship and of strategy more generally (Leiblein et al., 2018, Van den Steen, 2016), we take the view that accurate strategic foresight depends on whether the entrepreneur can anticipate the complexity in her strategy. There are two competing mechanisms through which experience can impact the accuracy of strategic foresight, depending on whether the experience addresses or exacerbates complexity. On one hand, experience increases the accuracy of strategic foresight if an entrepreneur learns about complexity that can apply to a future project. On the other hand, experience decreases the accuracy of strategic foresight if an entrepreneur learns about opportunities to augment her project, which introduces additional complexity. We argue that, when complexity increases rapidly across projects, the latter effect dominates the former. As a result, we theorize that as entrepreneurs gain experience across projects, they can introduce additional complexity that causes their strategic foresight to become less accurate.

In a study of 314 entrepreneurs across 722 crowdfunded hardware technology projects along with a program of qualitative interviews with serial crowdfunding entrepreneurs, we find that entrepreneurs make less accurate predictions as they gain experience across projects: they miss their predicted timeline to bring a product to market by a wider margin on each successive project, even as they actually give themselves more time on later projects. On average, an entrepreneur misses the timeline by a gap that grows by nearly six additional weeks on each subsequent project, and this effect persists: the gap between the predicted timeline and the actual delivery date continues to widen for later and later projects. We specifically study timeline predictions given that these predictions rely on strategic foresight and have meaningful strategic implications for customer value and firm survival, e.g., many entrepreneurs run out of money because they take more time than expected or, in other words, time

is money.

To explain this intriguing pattern, we show that, as entrepreneurs gain experience across projects, their future projects include additions that lead to more and more unforeseen interdependencies that they do not account for when making *ex ante* predictions. In addition to documenting these patterns in a quantitative analysis, our interviews provide a detailed view of how these mechanisms result in less accurate predictions. For example, one entrepreneur initially launched a Bluetooth LEGO brick for his customers to control motors and lights in their LEGO creations, e.g., a remote-controlled car. From this initial experience, the entrepreneur learned that it would be valuable for his next project to also add compatibility with LEGO sensors, e.g., the car could sense darkness to turn on a light. As the entrepreneur set the timeline for the subsequent product with more features, he did give himself more time than the previous product by setting the delivery date further out. However, he did not give himself enough time: we show that even with a simple addition, entrepreneurs encounter an increasing number of unforeseen interdependencies during implementation, suggesting an increase in complexity beyond the ability of an entrepreneur to foresee. Despite having a working prototype when making his prediction, this LEGO entrepreneur failed to foresee how adding this feature would have major consequences for many steps in the manufacturing process, like requiring different, more sophisticated tooling. His original manufacturer was no longer able to produce the product, and he went through seven different manufacturers before finding one who could produce the updated brick. Of course, he then missed his predicted delivery date.

This study makes three contributions. First, we outline the role of complexity in strategic foresight, proposing that complexity can serve as an alternate or at least more nuanced explanation for documented patterns of entrepreneurial failure and excess entry—characterized by prior literature as overconfidence—and that learning from experience may not be a cure-all solution to inaccurate strategic foresight. Second, we put forth the notion that strategic foresight comprises multiple interdependent predictions. Third, we argue that timeline predictions are strategically important with

direct implications for firm survival, and we provide suggestions for how managers can better predict timelines.

1.2 THEORETICAL BACKGROUND

1.2.1 COMPLEXITY AND STRATEGIC FORESIGHT

To unpack the potential effect of project experience on an entrepreneur’s strategic foresight, we need to first understand the role that complexity plays in this relationship.

Any given strategy that an entrepreneur might pursue, and need to make predictions about, entails complexity. By complexity, we mean the full set of interdependencies that exist between the components (or tasks) in the execution of a particular strategy (Simon, 1962). Thus, the complexity of a given strategy is a function of the number of components and the dependencies between those components, which together determine the total number of interdependencies that make up the full complexity of a strategy.¹

Complexity is not just an idiosyncratic characteristic of some strategies, but a core part of all strategies: across the board, recent efforts to formally define strategy specifically invoke complexity and interdependencies as first-order and necessary characteristics of what makes a course of action strategic at all (Nickerson & Argyres, 2018, Csaszar, 2018). Prior work emphasizes that the complexity of a strategy can itself be a source of competitive advantage that limits imitation (Rivkin, 2000), such that an entrepreneur could justify a strategy with high complexity despite its associated difficulty.

We take the view that entrepreneurship can be characterized as strategic foresight under complex-

¹To better align with our empirical context, our terminology for interdependencies differs subtly from how it is described in the NK modeling tradition (Levinthal, 1997, Kauffman, 1995). Although what we describe as the number of components or features roughly maps to N , what we refer to as the number of interdependencies is distinct from K . The general notion of K , as the *level* of interdependence, is traditionally defined as the dependencies that a single component $n \in N$ has on other components in the system. However, when we refer to the overall project complexity or the (total) number of interdependencies in a project, we mean the sum of all interdependencies across all components, which is closer to $N \times K$ rather than just K .

ity. Strategic foresight—and the ability to make accurate predictions related to a potential strategy—depends critically on the entrepreneur’s ability to anticipate the complexity and specific interdependencies she will face when later implementing or executing on the strategy (Gavetti & Menon, 2016). Entrepreneurs pursue more cognitively distant opportunities and then iterate on those opportunities as they learn (Gavetti, 2012). Due to the high velocity of entrepreneurial markets, entrepreneurship requires operating in novel settings of interdependencies (Eisenhardt & Bingham, 2017), where the entrepreneur, or anyone else for that matter, lacks prior experience with the interdependencies to be faced.

Thus, entrepreneurs have a particularly difficult challenge in anticipating the complexity they might face, limiting their effectiveness in making accurate strategic predictions.

1.2.2 EXPERIENCE AND STRATEGIC FORESIGHT: TWO CHANNELS

Our theory focuses on whether an entrepreneur can learn by experience across projects in such a way to improve strategic foresight for executing a project. By experience, we mean an entrepreneurial firm’s past exposure to the execution of tasks relevant to a given prediction. We identify two competing channels of learning through which past project experience might impact the accuracy of strategic foresight for a subsequent project.

LEARNING ABOUT PAST COMPLEXITY On one hand, an entrepreneur can learn about interdependencies by experiencing them when executing past projects: this experience would thus increase the accuracy of strategic foresight. Prior studies show that both organizations and individuals can learn from repeating interdependent tasks (Edmondson et al., 2007, Ethiraj et al., 2005). Denrell et al. (2004) show that learning in complex systems is best facilitated when there is continuity of personnel, like with serial entrepreneurs. In theory, if an entrepreneur and her organization execute the exact same project over and over again, she will have repeated instances of exposure to the interde-

dependencies inherent to that project because the full set of interdependencies that the entrepreneur must address for that project remains the same. With repetition, the entrepreneur should approach a full understanding of the system of interdependencies in the project. Improved knowledge of the interdependencies that she will face in execution leaves fewer interdependencies that she overlooks in her mental model when she makes predictions, enabling more accurate strategic foresight for future projects.²

While our theorizing here intentionally remains agnostic to heterogeneity in the performance of past experience, Section 1.5.1 leverages empirical findings to post hoc theorize that experience with underperformance in predictions for past projects may be beneficial for making a more accurate future prediction.

LEARNING ABOUT OPPORTUNITIES TO INCREASE COMPLEXITY On the other hand, project experience exposes an entrepreneur to opportunities to add new features to her next project, increasing its complexity: this experience risks decreasing the accuracy of strategic foresight. The entrepreneurship literature highlights how entrepreneurs identify opportunities to innovate in ways that emerge endogenously from their experience (Alvarez & Barney, 2007). These new opportunities are largely proximate to prior experience, involving incremental improvements to the prior pursuits.³ Acting

²Our theorizing here only focuses on the learning about interdependencies that comes about from actual experience with execution, independent of heterogeneity in the quality of that execution. The entrepreneur still needs to go through the motions of executing the project, which still exposes her to the interdependencies and gives her the knowledge she can take to future projects. We also consider an alternative and important behavioral mechanism of performance feedback, whereby past success or failure relative to aspiration levels may affect future behavior (e.g., Levinthal & March, 1993, Greve, 2003, 1998, Joseph & Gaba, 2015, Cho & Clough, 2015). Online Appendix Section A.1.9 further details this theoretical perspective and presents associated empirical tests evaluating the effect of experiencing success or failure on a past fundraising *campaign*—as opposed to our main focus of *project* execution—by raising more than or less than the desired amount of money from customers, respectively.

³Prior experience with project execution brings opportunities to an entrepreneur in two key ways. First, an entrepreneur discovers new opportunities when her experience exposes her to new information about customer needs and ways to serve the market. Second, an entrepreneur creates new opportunities through an enactment process where, in the course of prior experience, she may devise new ways of combining preexisting knowledge.

on these new opportunities by making even just incremental additions to the product increases the complexity by adding new, previously unencountered interdependencies (Anderson, 1999). Thus, to make accurate predictions about the opportunity, the entrepreneur would have to be able to account for those interdependencies. In this way, gaining new knowledge through experience could even exacerbate the challenge of complexity and, as a result, increase the number of ways strategic foresight could be inaccurate (Townsend et al., 2018a). To conceptually pinpoint the net effect of project experience on strategic foresight, we now need to identify which of these two channels dominates.

1.2.3 DOMINANCE OF INCREASING COMPLEXITY

We contend that—under certain conditions common to entrepreneurial settings—new complexity can outweigh the benefits of experience. The argument follows from assumptions we can make about the shape of the *Project Complexity* curve and the *Learning* curve, described here and visually illustrated in Figure 3.3. On the one hand, as traced by the increasing *Learning* curve, as an entrepreneur gains experience and learns she can anticipate an increasing number of *Foreseen Interdependencies*. On the other hand, as an entrepreneur gains experience across projects, she also learns about opportunities to add features to expand her next product. Adding these new features increases the *Project Complexity* by adding new, previously unencountered interdependencies. We argue that the latter effect can dominate the former: when the *Project Complexity* curve increases faster than the *Learning* curve, the entrepreneur ultimately faces an increasing number of *Unforeseen Interdependencies* that will be overlooked in the prediction process and impair strategic foresight.

Under the assumptions detailed below—at least a linearly increasing project complexity curve and a concave learning curve—we theorize that past project experience has a negative relationship with the accuracy of strategic foresight for a subsequent project.

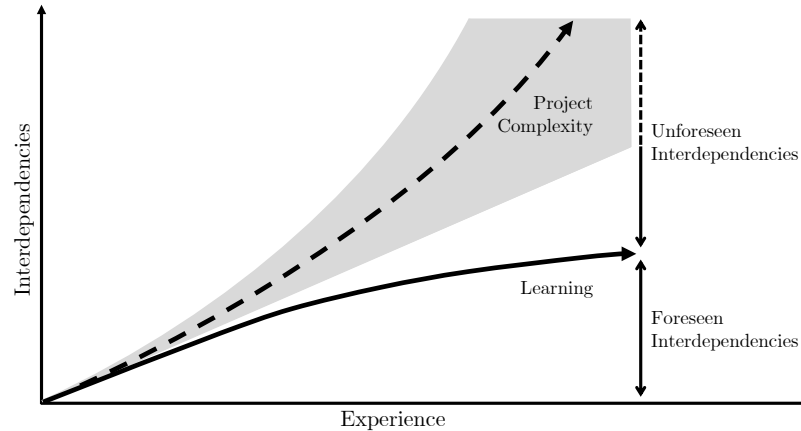


Figure 1.1: Conceptual Model of Experience and Interdependencies. The vertical axis *Interdependencies* represents the number of interdependencies. The horizontal axis *Experience* represents an entrepreneur’s level of project execution experience. The *Project Complexity* curve illustrates the total number of interdependencies in projects pursued by an entrepreneur at different levels of *Experience*. The shaded grey area reflects the range of possible *Project Complexity* curves: linear as a lower bound and geometric as an upper bound; the black dotted line illustrates one possible scenario. The *Learning* curve illustrates the total number of interdependencies foreseen (*Foreseen Interdependencies*) by an entrepreneur at different levels of *Experience*. As *Experience* increases, the gap between the *Learning* curve and the *Project Complexity* curve increases and, as a result, the ratio of *Unforeseen Interdependencies* to *Foreseen Interdependencies* also increases.

INCREASING PROJECT COMPLEXITY Even when an entrepreneur makes merely incremental additions to a previous project, complexity increases. Adding a new feature requires adding one or more tasks interdependent with some or many tasks in their system of activities (Ethiraj et al., 2012).⁴ As a result, each new feature added must increase the total number of interdependencies in the project. The overall theoretical argument follows from the minimum baseline premise that the total number of interdependencies increases at least linearly, which assumes that the entrepreneur would have to add at least one component or task in a subsequent project and that the addition should be at least as interdependent as other components that already exist in the prior project.

That said, we posit that in most entrepreneurial ventures complexity can increase faster than lin-

⁴Our study focuses on hardware technology projects that force the entrepreneur to integrate components at some level: if there were no interdependencies, there would be no opportunity for value creation by the entrepreneur as the raw inputs could just be purchased separately by customers with no loss of value.

early, well above the minimum assumption needed for the theory to hold. First, it can be the case that an entrepreneur adds multiple features or tasks in a subsequent project, particularly for a nascent entrepreneur improving on a sparse project far from a dominant design. Second, for projects with highly interdependent components, the addition of a single component can lead to a faster-than-linear increase in the total number of interdependencies; at the extreme, the number of interdependencies can increase geometrically.⁵ While both these conditions vary based on context, entrepreneurs engaged in launching a new product—particularly a new hardware technology as in our empirical context—likely meet both of these conditions. The next section describes how these assumptions hold in context.

Product development entails highly interdependent components and tasks (Ulrich et al., 2020), meaning that the entrepreneur faces a complex system that is inherently nonlinear (Anderson, 1999, Townsend et al., 2018a). As a result, adding new components leads to a cascade of new interdependencies which grows rapidly and may outpace the comparatively incremental discovery of interdependencies encountered in past experience. Thus, as an entrepreneur gains experience across projects and implements new features for a subsequent project, demonstrating strategic foresight requires that the entrepreneur navigate more complexity, and perhaps substantially more, than previously faced.

BOUNDING LEARNING As the entrepreneur takes on more complexity, the potential benefits of learning about interdependencies from prior experience are increasingly limited. Entrepreneurs operating in complex systems rely on simplified mental models that only account for a subset of the total interdependencies. Describing this simplification process, Eisenhardt & Bingham (2017) detail entrepreneurs' use of simple models, Csaszar (2018) compares different simplified representations of complexity, and Gavetti (2012) outlines the necessity of associative thinking. Although the frameworks proposed in these studies make some distinctions, the broad consensus is that entrepreneurs simplify the system of interdependencies in making judgments. By definition, these simplified mod-

⁵For instance, a project with X components that are all interdependent with one another would have $X(X-1)$ total interdependencies, a function geometrically increasing in X .

els are incomplete. Furthermore, due to the cognitive constraints on the number of interdependencies an entrepreneur is able to consider (Simon, 1990, 1969), these models will be less complete in more complex systems. As entrepreneurs implement increasingly complex successive projects, the portion of the total interdependencies the entrepreneur is able to foresee decreases. Consequently, expanding the total number of interdependencies increasingly penalizes the accuracy of an entrepreneur's strategic foresight.

1.3 HYPOTHESIS DEVELOPMENT

From this conceptual viewpoint, we now develop a series of hypotheses situated in our empirical context: hardware technology entrepreneurs engaging in product crowdfunding on Kickstarter. In particular, we consider entrepreneurs serially crowdfunding across multiple distinct projects of the same subtype. In order to validate our aggregate empirical patterns and understand potential micro-mechanisms, we conduct a program of qualitative interviews with 11 entrepreneurs from our sample.⁶ We weave in qualitative findings from these interviews into our hypothesis development purely for context and clarity. These examples and anecdotes are not intended as empirical proof for the theory, but as transparent illustrations of the logic underlying the theoretical mechanisms.⁷

1.3.1 INCREASING UNFORESEEN INTERDEPENDENCIES

We predict that as entrepreneurs gain experience across projects, they suffer from an increasing number of unforeseen interdependencies. Each time an entrepreneur executes a project, she gains experience designing, prototyping, manufacturing, and delivering a product. Consider an entrepreneur repeatedly executing the exact same project with the same set of tasks and interdependencies again and

⁶Online Appendix Section A.1.1 describes our qualitative interview process.

⁷We intend for the theory and hypotheses to arise from conventional deductive arguments, which we then ground in our specific context using examples and quotes from the qualitative interviews, rather than using the qualitative evidence as a basis for inductive theory development.

again. We would expect her to learn and update her prior beliefs about the set of interdependencies for the next related pursuit (Raveendran et al., 2020). Under this scenario, experience improves strategic foresight. Consider the case of MaskCo, which creates sound-reactive LED masks. In 2015, MaskCo launched its first mask project on Kickstarter: a jaguar design outlined by basic LED strips. On its initial project, MaskCo experienced unforeseen manufacturing challenges, leading to production delays. However, if MaskCo continues to produce this exact same mask again and again, we would not expect it to continue to suffer from the same unexpected manufacturing challenges. Rather, we would expect the number of unforeseen interdependencies to decrease.

However, this ceases to be the case when an entrepreneur implements new features discovered while implementing past projects. In this scenario, the total number of interdependencies increases relative to the previous project. MaskCo's initial Kickstarter experience exposed it to additional opportunities to innovate based on consumer feedback suggesting demand for additional design options, leading MaskCo to add a host of new design options—including an owl, wolf, fox, skull, robot, wildcat, and even a version with President Obama's face—on its subsequent project. The MaskCo entrepreneur also discovered new ways of combining pre-existing knowledge: the initial mask would only light up in response to sound, but the entrepreneur deduced that it would be valuable to have pre-programmed light patterns so the mask could also light up without sound.

Each new added feature interacts with some or many of the tasks and components required to complete the prior project, thus introducing new interdependencies. But when an entrepreneur makes predictions about this more complex product specification and the timeline on which she will deliver it, our theory suggests that the entrepreneur may do so with an incomplete view of the interdependencies that might arise. If new features added to a subsequent project increase the total number of project interdependencies in excess of the foreseen interdependencies gained through learning on prior projects, the number of unforeseen interdependencies will increase on each subsequent project.

Hypothesis 1 *As entrepreneurs gain experience from past projects, they encounter an increasing number*

of unforeseen interdependencies on their next project.

To illustrate this hypothesis in context, we continue with the example of MaskCo and highlight the seemingly small choice to introduce packaging to the company's subsequent project. Adding packaging to a product that is even otherwise the same introduces significant complexity given all the ways the new packaging is interdependent with the existing production tasks. This addition required MaskCo to arrange for the packaging to occur at a separate plant, which necessitated coordinating shipping between the plants and hiring a contractor to facilitate communication in a different language between the manufacturer and the packaging plant. Then, when the quality of the first finished batch was poor, correcting the problem took even more time given the additional interdependency of the finished product with packaging. Going back through the whole process to correct the problem and then repackage the products cost MaskCo an additional month. Then, the new packaging meant that the finished products could no longer be shipped by the shipping company used previously, so MaskCo ultimately had to move all the stock to a different warehouse for shipment.

Our interviewees repeatedly emphasized unexpected organizational issues that came up during execution. Given Kickstarter's requirement to have a working prototype before fundraising, many if not most of the interdependencies intrinsic to the product itself were already known prior to launching the project. However, "the prototypes are all hand made—they're more of a unique product that has more time put into it—but when you're doing production, you're not spending that much time on every single unit. You're doing large volume. That's where we end up having problems" (GPSCo CEO). Another entrepreneur shared, "our [second product] was more complicated because organizing all the different sourcing was a lot more difficult. For [the first product], it was basically, 'go to one supplier and then just put in an order.' But with [the second product], there was a lot of back and forth with a bunch of different suppliers" (CircuitsCo CEO). Indeed, the MaskCo entrepreneur noted that the ultimate set of steps required to add packaging involved "things [he] never thought about" in working with other organizations.

1.3.2 STRATEGIC FORESIGHT AS MULTIPLE PREDICTIONS

We now turn to how entrepreneurs respond when they encounter unforeseen interdependencies that conflict with the strategic foresight of their initial predictions. As a starting premise, we characterize strategic foresight as a set of multiple predictions. When our theory suggests that entrepreneurs make increasingly inaccurate or infeasible predictions on each subsequent project, we mean that with respect to the aggregate of all the entrepreneur's predictions that comprise their strategic foresight as a whole. The individual predictions are fundamentally connected: entrepreneurs have the choice to absorb the inaccuracy in one prediction while satisfying another prediction.

It is important to discuss predictions in context because strategic foresight in different contexts comprises different dimensions on which entrepreneurs make predictions. Crowdfunding entrepreneurs make two important, and readily observable, predictions: product specification and delivery timeline, meaning the date they will deliver the product to customers. Entrepreneurs make these predictions publicly to prospective customers who finance a project on the possibility that they will receive the specified product by the specified date. Based on our qualitative interviews, we find that entrepreneurs make these predictions first by detailing an anticipated product specification, and then setting a delivery date by breaking the production process down into concrete interdependent tasks, predicting the timeline for each task, and aggregating those timelines. In most cases, entrepreneurs also try to be conservative by adding some buffer time to their overall timeline.⁸

Product specification and delivery timeline are connected in such a way that the prediction relative to one can be met at the expense of the other. For example, if an entrepreneur makes an inaccurate timeline prediction, she could still choose to meet the timeline prediction by delivering a product that fails to meet the product specification (and vice versa).⁹ In principle, an entrepreneur could choose to

⁸Online Appendix Section A.1.3 elaborates on this prediction process for hardware technology projects and provides qualitative context from entrepreneur interviews.

⁹Never delivering a product would be an asymptotic combination of these two ways of missing a prediction, i.e., delivering a product of zero value with an infinite delay. We exclude this situation from our empirical

prioritize a predicted timeline by allocating a fixed amount of time to a project, even if the predicted product specification is not fully achieved, so she can move on to other activities.

1.3.3 PRIORITIZING PRODUCT SPECIFICATION OVER TIMELINE

However, we argue that entrepreneurs in the crowdfunding context—and perhaps many in other settings—prioritize achieving the predicted project specification rather than adhering to the initially predicted delivery date. In other words, given inaccurate strategic foresight, most entrepreneurs tend to continue working towards achieving a predicted product specification, even if it requires going beyond the originally predicted delivery date. This tendency to prioritize achieving product specification over meeting a timeline follows if an entrepreneur holds certain beliefs about customer preferences and the resulting consequences of achieving (or not) either predicted dimension. While there are meaningful consequences for delay,¹⁰ these consequences are overshadowed by both the negative consequences of failing to deliver the specified product as well as the positive benefits of succeeding in doing so. If a customer receives a product below the promised specification, this can cause severe reputational damage to the entrepreneur. However, delivering a product as specified (even a delayed product) can still lead to brand-building testimonials and organic growth. Additionally, succeeding in delivering the specified product allows an entrepreneur to get feedback on her actual intended product specification which she can then use to develop future projects.

When inaccurate strategic foresight leads to unforeseen interdependencies that make it impractical to achieve both initial predictions, entrepreneurs can choose which prediction they will ultimately prioritize and achieve and which to relegate and fail to address. We argue that most entrepreneurs prioritize achieving the predicted product specification over the predicted delivery date. As a result, as entrepreneurs gain experience implementing projects and encounter an increasing number of un-

analysis because this situation is rare and some potentially substantial number of those situations involve fraud by the entrepreneur (Mollick, 2015).

¹⁰Online Appendix Section A.1.1.1 expounds and quantifies these consequences of delay.

foreseen interdependencies on subsequent projects, requiring additional effort beyond what was predicted (Ethiraj, 2007), we expect achieving their predicted product specification requires failing to achieve their predicted timeline by increasing margins. This will manifest in increasing delays.

Hypothesis 2 *As entrepreneurs gain experience from past projects, they fail to achieve their predicted delivery date on their next project by a wider margin.*

Without exception, our interviews with crowdfunding entrepreneurs confirm this tendency to achieve their predicted product specification at the expense of their predicted delivery date. One explained, “At the end of the day, you have to make the decision: Do I want to ship a product that we don’t feel meets the needs of the customer just to be able to ship it and be done with it? Or do we want to delay and end up shipping a quality product? I always want to ship a quality product” (GPSCo CEO). Another entrepreneur believed that “consumers can delay gratification for something better” (TabletCo CEO). To put it another way, “We wanted to first be able to deliver the highest-quality parts we could, and then second to do as best we can to deliver them on time” (3DPrintCo CEO).

We observe this tendency in both the LEGO brick and MaskCo entrepreneurs mentioned previously. The LEGO brick entrepreneur referenced in the introduction could have decided to deliver a product on the predicted delivery date that did not perform the predicted function of interfacing with LEGO sensors. Similarly, the MaskCo entrepreneur could have delivered a mask in whatever state it was in (perhaps without packaging) by the predicted delivery date. However, both entrepreneurs chose to delay in order to continue striving to meet the predicted product specification. The discussion highlights other prominent examples where entrepreneurs—like Elon Musk of Tesla—exhibit this tendency to spend more time working towards their predicted product specification rather than adjusting their product specification to meet the predicted allocation of time resources.

1.4 EMPIRICAL METHODS

1.4.1 CONTEXT

In order to test these hypotheses, we need a sample of entrepreneurs who complete multiple projects over time with clearly defined markers for experience, complexity, predictions, and outcomes. The crowdfunding platform Kickstarter provides an ideal setting that meets these criteria. Kickstarter, founded in 2009, is a popular crowdfunding platform that connects entrepreneurs to customers. Customers pre-purchase specific products that the entrepreneurs promise to deliver by a future date. This fundraising process requires Kickstarter entrepreneurs to provide several predictions, including the features and qualities of the product they will produce and the timeline on which they will deliver the product. This setting allows us to identify metrics to capture each of the characteristics and outcomes of interest outlined in our hypotheses. Figure 2.2 provides specific examples of these metrics using the series of projects implemented by one of the entrepreneurs in our sample.

Using Kickstarter projects favorably standardizes several characteristics. All hardware technology projects are required to have a working prototype before they can raise capital, helping to reduce some of the variation in the starting point of new projects (Kickstarter PBC, 2020). The crowdfunded capital then funds the manufacturing and distribution of the product at scale. In addition, the platform is all-or-nothing, meaning that if the project does not reach the target financing level, the pledges are refunded to the customers and the entrepreneur does not receive any capital. As a result, we can assume that the entrepreneurs have sufficient financial resources to deliver the product relative to their expectations.

Although some associate Kickstarter with fun trinkets and games, our study focuses on manufactured hardware technology, the most complex products on Kickstarter and among the most complex that an entrepreneur could generally pursue.¹¹ First, the value of these products hinges on precisely

¹¹Online Appendix Section A.1.2 details the high and increasing degree of complexity in crowdfunded hard-



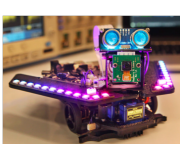
Name & Date	Image	Features Rank	Delay Duration	Delay Duration/ Predicted Time	Unforeseen Interdependencies
Ringo Feb 25, 2015		1 (Fewest)	66 Days	92%	"Our machine refused to pick up the programming ports...this programming port was just a bit [too] heavy."
Wink Oct 28, 2015		2 (Middling)	73 Days	115%	"The testing procedure is taking longer than expected...finding a few units with bad motors." "We ran out of motors and our replenishment shipment was held up."
Spirit Rover Sept 28, 2016		3 (Most)	340 Days	254%	"Found two mistakes on the boards... fixed with an extra step on our end, but I should have known better on both of these." "We finally found sources for all the screws, fasteners, washers, nuts, and spacers. I was surprised and unprepared at how difficult this part was going to be." "I made a mistake with two of the cables... as they are too short."

Figure 1.2: Example Products by Entrepreneur Over Time. All projects by Plum Geek Robotics, founded by Kevin King, in the robotics subtype of the technology category. The Unforeseen Interdependencies column provides quotations from updates by the entrepreneur. All other variables mirror those defined in the paper.

and accurately addressing a large number of interdependencies. If a wire is cut a nanometer too short, it may not connect the necessary circuits for the product to function. In contrast, if the pair of dice in a board game is produced a nanometer smaller than planned, it has virtually no impact on the other game pieces. Second, modern manufacturing requires an international supply chain with multiple suppliers from different organizations, e.g., distinct suppliers for all the parts, assembly, packaging, and international shipping.¹²

1.4.2 DATA AND SAMPLE

We construct a sample of Kickstarter entrepreneurs who complete multiple projects of the same project subtype. This should, in principle, keep experience gained on a past project relevant to the next ware technology products.

¹²Online Appendix Section A.1.3 further expounds the complexity inherent in this context as well as the implications of that complexity for the prediction process.

project, which is ideal for reaping the benefits of learning. We collect basic project data and characteristics for all Kickstarter projects from Web Robots, which runs a monthly scrape of all past and present Kickstarter projects. We identify the 394 entrepreneurs with two or more projects that met the fundraising goal in one of the main project subtypes in the hardware technology space (i.e., gadgets, 3D printing, hardware, camera equipment, sound, DIY electronics, wearables, robots, and fabrication tools) with predicted delivery dates prior to the date of our analysis. We look specifically at entrepreneurs with multiple projects that meet the fundraising target because they gain execution experience from actually having to produce and deliver these projects. In order to maximize the potential impact of learning, we further segment our sample to the entrepreneurs who specialize in one of the selected project subtypes, refining our sample to 326 entrepreneurs.¹³ After reviewing each entrepreneur’s profile, we also exclude 12 entrepreneurs whose circumstances are disqualifying (e.g., a large, established company launches the campaign) or where it is apparent we have incomplete data (e.g., the entrepreneur is clearly doing many other projects outside of Kickstarter, in which case our data set does not capture much of their relevant experience).

These criteria result in a final sample of 314 entrepreneurs who created 722 projects from September 2010 through June 2019. For each of these projects, we scrape comprehensive information from its Kickstarter pages, including the most recent 100 comments and all project updates posted by the entrepreneur. We manually collect data on actual delivery time and number of features. We link Kickstarter entrepreneurs with their Crunchbase profiles to track their external funding over time.¹⁴

¹³ While an entrepreneur could intentionally shift to a product subtype “distant” from her prior experience (e.g., [Eggers & Song, 2015](#)), this possibility falls outside the scope of this study.

¹⁴ Online Appendix Section A.1.4 provides additional detail about the data collection and aggregation process.

1.4.3 VARIABLES

DEPENDENT VARIABLES: FEATURES AND UNFORESEEN INTERDEPENDENCIES

As a starting point, we define a set of measures to test the basic assumption leading into Hypothesis 6: entrepreneurs pursue increasingly complex projects, i.e., projects with greater total interdependencies. An ideal measure would exactly measure the total interdependencies in a predicted project, but this is impossible to identify based on the public information available since we cannot see inside the product or organization. Instead, we identify a product's level of features relative to the other product(s) by the same entrepreneur. We hired five independent reviewers to rank each entrepreneur's set of products by number of features and then aggregated the rankings for each product across reviewers.¹⁵ Specifically, we hired individuals with relevant educational and professional experience in computer programming, mechanical engineering, and robotics. The following measures are intended to at least roughly correlate positively with the total interdependencies in a planned project.

FEATURES MOST *Features Most* is a binary indicator equal to 1 if the product has the most features compared to the other products by the same entrepreneur (and 0 otherwise).

FEATURES RANK *Features Rank* is the relative rank of the product compared to other products by the same entrepreneur, e.g., if an entrepreneur completed two products, the product with the fewest features would have a *Features Rank* of 1 and the product with the most features would have a *Features Rank* of 2.

FEATURES PERCENTILE *Features Percentile* specifies the relative percentile of a project for an entrepreneur, e.g., if an entrepreneur had three projects with *Features Rank* equal to 1, 2, and 3, the

¹⁵Online Appendix Section A.1.4 provides additional details on the background of each reviewer as well as the ranking and aggregation process.

corresponding *Features Percentile* would be 0%, 50%, and 100%, respectively.

UNFORESEEN INTERDEPENDENCIES We then construct a measure of unforeseen interdependencies in a direct test of Hypothesis 6. *Unforeseen Interdependencies* is the total number of updates posted by the entrepreneur during project execution—after the fundraising campaign has ended and before the product is delivered—that cite unforeseen interdependencies. A member of our research team reviewed the most common words contained in updates relevant to unforeseen interdependencies. They identified two categories of relevant words. The first set relate to issues being unforeseen, which include the words (or any variants): unforeseen, unexpected, and unanticipated. The second set relate to typical interdependence-related issues that come up in our context, which include the words (or any variants): manufacturing, production, assembly, and factory. When defining *Unforeseen Interdependencies*, we include all updates that contain words from either set.¹⁶

DEPENDENT VARIABLES: DELIVERY TIME METRICS

DELAY INDICATOR If the actual delivery date is after the predicted delivery date or if the project has not yet shipped and the predicted delivery date is prior to the date when we collected our sample, *Delay Indicator* is set equal to 1 (and 0 otherwise). We identify *Delay Indicator* for 95% of projects in our sample.

DELAY DURATION *Delay Duration* is the time (in days) between the actual delivery date and the predicted delivery date. We identify the delay for 89% of our sample; for comparison, Mollick (2014) identifies outcomes for 81% of his sample.

¹⁶Online Appendix Section A.1.6 shows that all results hold if we define *Unforeseen Interdependencies* to contain updates with words from both sets.

PREDICTED TIME To test whether *Delay Duration* is driven by more aggressive predictions versus missing static predictions by wider margins, we define *Predicted Time* as the time (in days) between the end of the fundraising campaign and the predicted delivery date.

ACTUAL TIME *Actual Time* is the time (in days) between the end of the fundraising campaign and the actual delivery date, i.e., the sum of *Predicted Time* and *Delay Duration*.

MAIN INDEPENDENT VARIABLE: PROJECT EXPERIENCE

The main independent variable *Project Experience* measures an entrepreneur's total execution experience as her number of projects prior to her current project and of the same subtype. We only count projects that meet the funding threshold because they provide the entrepreneur with execution experience that exposes her to project interdependencies.

CONTROL VARIABLES

Entrepreneur fixed effects control for any time-invariant variation among entrepreneurs in our sample, so we add additional controls for other types of entrepreneur experience characteristics that may change over time, as well as project-specific characteristics.

We control for other types of entrepreneurial experience with executed *projects* (that meet the funding threshold) and attempts at funding *campaigns* (most of which become projects). Given the potential impact of fundraising failure on behavior,¹⁷ *Failed Campaign Experience* is the cumulative count of Kickstarter campaigns of the same product subtype conducted by the entrepreneur where

¹⁷Compared to the average 70.9% failure rate of technology Kickstarter projects, only 10.5% of the campaigns attempted by the entrepreneurs in our sample failed to reach their funding threshold. This is likely driven by key differences between the serial-project entrepreneurs in our sample who generally treat their projects as full-time jobs and the average person who casually launches a project more as a hobby.

those campaigns did not reach their funding threshold.¹⁸ In a similar vein, to account for the degree and direction of deviation from the funding threshold, *Prior Campaign Funding Deviation* is the percentage by which the entrepreneur's prior campaign exceeded (or missed) its funding threshold. Another way past performance could impact an entrepreneur's behavior on subsequent projects is the number of days by which the entrepreneur missed (or beat) their predicted timeline on the past project. *Prior Project Delay* is the entrepreneur's prior project's *Delay Duration* divided by *Predicted Time*.

We also include controls for changes in the entrepreneur's circumstances over time. Simultaneous execution of multiple projects could impact performance as compared to projects that are the sole focus of an entrepreneur. *Execution Overlap* is a binary indicator equal to 1 if the execution start date of the current project comes before the execution completion date of the prior project (and 0 otherwise).¹⁹ To control for changes in entrepreneur quality over time, we use Crunchbase data and define *External Funding* as a binary indicator of whether the entrepreneur had raised venture capital funding prior to launching the current project. To account for the impact of switching industries documented by Eggers & Song (2015), we define *New Category* as a time-variant binary indicator of whether the project immediately prior to the focal project was of a different subtype. We also control for general experience and learning that may accrue to the entrepreneur naturally over time and separate from project execution, with *Elapsed Time* defined as the number of days since the entrepreneur launched her first successful project of the same subtype as the current project. *Baseline Updates* is the total number of updates posted prior to the end of the fundraising campaign, which allows us to control for the entrepreneur's time-variant propensity to post updates across projects.

¹⁸ *Project Experience* plus *Failed Campaign Experience* is the total number of Kickstarter campaigns of the specific product subtype that the entrepreneur had launched; including both of these variables together also controls for the total number of campaigns in aggregate, which would be collinear.

¹⁹ We look at overlap in execution rather than fundraising given that executing a project takes substantial time and other resources. This overlap only occurs in 4.7% of our sample (34 out of the 722 projects). This makes sense given that the entrepreneurs interviewed noted that executing even a single project is generally a full-time job and the ideas for subsequent projects come through executing past projects.

We also control for project characteristics determined ex ante to initiating the fundraising campaign. These variables control for whether heterogeneity in project characteristics account for heterogeneity in measured outcomes. *Funding Period* is the time (in days) that the project accepted contributions; this window is set before the fundraising campaign launches and cannot be changed after the fact. *Funding Reward Tiers* is the total number of rewards available for funding backers to purchase. *Funding Reward Size* is the median price of the rewards available for funding backers to purchase. *Funding Threshold* is the amount of money (in thousands of USD) the entrepreneur set out to raise; since this amount is set at the start of the *campaign* and cannot be adjusted, all *projects* meet or exceed this threshold.

In addition, we control for project characteristics determined ex post after the fundraising period. We include these ex post controls in regressions where the dependent variable is realized after the fundraising period. *Funding Exceeded* is the amount of money (in thousands of USD) the project raised in excess of the *Funding Threshold*; Mollick (2014) finds that the degree to which projects exceed the funding threshold associates with delay. *Funding Backers* is the total number of people (in thousands) who contributed to the project.

1.4.4 DESCRIPTIVE STATISTICS

Table 1.1 provides summary statistics.²⁰ To validate our measures, we compare our sample of 722 technology projects to the 843 technology projects in Mollick (2014): our sample has an average *Funding Threshold* of \$23,272 (versus \$21,177) and *Funding Period* of 33.34 days (versus 40.28 days). In addition, Mollick (2014) uses a similar manual process to collect actual delivery dates and finds that “only 24.9% of projects delivered on time” (or 75.1% of projects are delayed). Our sample identifies a similar pattern, where 76.3% of projects are delayed.

Looking at the pairwise correlations between each of our independent variables in Table 2.2, we

²⁰Online Appendix Section A.1.5 provides additional statistics and visualizations of variable distributions.

Table 1.1: Summary Statistics. 722 project-level observations. *Actual Time* and *Delay Duration* are based on 644 observations, and *Delay Indicator* is based on 686 observations. *Funding Threshold* (USD), *Funding Exceeded* (USD), and *Funding Backers* (count) are all in thousands.

Dependent Variables	Mean	Std. Dev.	Min	Max
Features Most	0.45	0.50	0	1
Features Rank	1.68	0.74	1	6
Features Percentile	0.50	0.48	0	1
Unforeseen Interdependencies	3.39	4.11	0	31
Delay Indicator	0.76	0.43	0	1
Delay Duration	70.72	114.84	-77	946.60
Predicted Time	90.48	52.97	5	414
Actual Time	159.25	142.81	10	1, 231.60
Independent Variables	Mean	Std. Dev.	Min	Max
Project Experience	0.70	0.74	0	5
Failed Campaign Experience	0.10	0.36	0	4
Prior Campaign Funding Deviation	3.24	8.86	-1	86
Prior Project Delay	0.42	1.02	-1	11
Execution Overlap	0.05	0.21	0	1
External Financing	0.09	0.29	0	1
New Category	0.03	0.17	0	1
Elapsed Time	322.72	426.76	0	2, 458
Baseline Updates	6.73	5.12	0	40
Funding Period	33.34	10.02	2	60
Funding Reward Tiers	9.72	5.07	1	34
Funding Reward Size	234.83	493.38	4	5, 995
Funding Threshold	23.27	31.40	0.02	261.96
Funding Exceeded	102.54	273.62	0	3, 351.36
Funding Backers	0.95	2.19	0.001	28.14

note the expected correlation (0.748) between *Funding Exceeded* and *Funding Backers*, since each new backer contributes additional funds to the project. We re-run all regressions taking turns excluding each of these variables and do not observe any meaningful changes to the results. In addition, there is an expected correlation (0.697) between *Project Experience* and *Elapsed Time*, given that each subsequent project occurs at a later time. All the results hold if we remove *Elapsed Time* from the regressions.

Table 1.2: Pairwise Correlation Matrix of Independent Variables.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
(1) Project Experience	1														
(2) Failed Campaign Experience	0.173	1													
(3) Prior Campaign Funding Deviation	0.251	-0.059	1												
(4) Prior Project Delay	0.306	0.009	0.162	1											
(5) Execution Overlap	0.161	0.027	0.008	0.129	1										
(6) External Financing	-0.004	-0.091	0.124	0.036	-0.026	1									
(7) New Category	-0.036	-0.006	0.063	0.001	-0.001	-0.001	1								
(8) Elapsed Time	0.697	0.064	0.241	0.334	-0.049	0.124	0.001	1							
(9) Baseline Updates	-0.172	-0.156	-0.081	-0.038	-0.113	0.081	-0.027	-0.086	1						
(10) Funding Period	-0.010	-0.121	0.093	0.012	-0.121	0.061	-0.039	0.033	0.193	1					
(11) Funding Reward Tiers	-0.056	-0.163	0.026	-0.004	-0.122	0.062	-0.046	0.024	0.277	0.218	1				
(12) Funding Reward Size	0.005	-0.068	0.008	0.035	-0.019	-0.007	-0.054	0.057	0.148	0.010	0.017	1			
(13) Funding Threshold	-0.084	-0.160	0.002	0.050	-0.081	0.190	-0.052	0.065	0.215	0.156	0.187	0.254	1		
(14) Funding Exceeded	-0.076	-0.092	0.175	-0.003	-0.070	0.247	-0.054	-0.005	0.181	0.138	0.140	0.091	0.307	1	
(15) Funding Backers	-0.093	-0.091	0.110	0.004	-0.070	0.276	-0.053	-0.051	0.149	0.152	0.161	-0.085	0.227	0.748	1

Given their importance, Figure 1.3 visualizes the distributions of *Predicted Time* and *Actual Time*. The distribution of *Actual Time* is shifted and skewed to the right of the distribution of *Predicted Time*, of course because the vast majority of projects are delayed.

1.4.5 STATISTICAL MODEL

We estimate ordinary least squares (OLS) models across all analyses. These models include fixed effects to control for several dimensions of otherwise unobserved heterogeneity that could correlate with the observed independent variables. Entrepreneur fixed effects control for time-invariant entrepreneur characteristics, such as natural talent, intelligence, work ethic, etc. Product subtype fixed effects absorb any heterogeneity between the various categories of projects, e.g., DIY electronics versus 3D printing. Month fixed effects control for seasonal cycles, e.g., if projects that predict delivery




Figure 1.3: Density Plot of *Actual Time* and *Predicted Time*. The distribution of *Actual Time* is shifted and skewed to the right compared to the distribution of *Predicted Time*. We adjust the bandwidth to smooth the distributions.

dates in December are more likely to delay due to the holidays, month fixed effects would account for that seasonal heterogeneity. Year fixed effects control for any factors that change year to year but are common to all entrepreneurs who launch new projects in a given year. To account for potential correlation in the error term across projects by the same entrepreneur, we cluster robust standard errors at the entrepreneur level.

The models using the dependent variables *Features Most*, *Features Rank*, *Features Percentile*, and *Predicted Time*—determined ex ante to launching the fundraising campaign—include only the controls for project characteristics that exist ex ante and exclude the control variables realized ex post, *Funding Exceeded* and *Funding Backers*.

1.5 RESULTS

Hypothesis 6 predicts that, as entrepreneurs gain experience, they encounter an increasing number of unforeseen interdependencies. Before we look at this directly, we first validate a key assumption leading to this hypothesis: entrepreneurs make their product specification more complex as they gain experience. We examine this by looking at the relationship between *Project Experience* and three measures of how complicated the proposed product specification is in terms of its observable features. In the first three columns of Table 3.3, we find that *Features Most* ($p = 0.047$), *Features Rank* ($p = 0.000$), and *Features Percentile* ($p = 0.016$) are all positively related to *Project Experience*.²¹ Each subsequent project is 11.7% more likely to be the highest-ranked project in terms of number of features. The ranking of each subsequent project increases by an average of 0.37 in absolute terms or 13.8% on a percentile basis. If additional features increase the number of interdependencies, we posit that more experienced entrepreneurs take on projects with more total interdependencies.

To explicitly test Hypothesis 6, we examine the effect of experience on the number of unforeseen interdependencies. In column 3 of Table 3.3, we find that *Unforeseen Interdependencies* ($p \sim 0.000$) is positively related to *Project Experience*. On each subsequent project, entrepreneurs disclose encountering 1.3 additional unforeseen interdependencies. This increase in unforeseen interdependencies is consistent with our theory of decreasing prediction accuracy in increasingly complex systems when bounded rationality limits the learning that might attenuate unforeseen interdependencies. *Ln Unforeseen Interdependencies* ($p = 0.002$) also positively associates with *Project Experience*. Each subsequent project increases unforeseen interdependencies by 21.0%.²²

Hypothesis 7 predicts what entrepreneurs will do when they make inaccurate predictions. Specif-

²¹For the binary indicator variable *Features Most*, the results hold when using a conditional fixed-effects logit model.

²²Online Appendix Section A.1.7 shows there is a significant and positive relationship between increasing unforeseen interdependencies and increasing number of features as well as between delay and increasing number of features.

Table 1.3: Features and Unforeseen Interdependencies. Ordinary least squares (OLS) estimation. Robust standard errors are clustered at the entrepreneur level. p -values are shown in parentheses.

	Features Most	Features Rank	Features Percentile	Unforeseen Interdependencies	Ln Unforeseen Interdependencies
Project Experience	0.117 (0.047)	0.370 (0.000)	0.138 (0.016)	1.298 (0.000)	0.205 (0.002)
Failed Campaign Experience	0.075 (0.506)	0.169 (0.245)	0.081 (0.401)	0.672 (0.066)	0.185 (0.057)
Prior Campaign Funding Deviation	0.001 (0.779)	0.000 (0.932)	0.001 (0.858)	0.005 (0.772)	-0.000 (0.908)
Prior Project Delay	0.011 (0.758)	-0.034 (0.535)	0.003 (0.941)	-0.698 (0.000)	-0.128 (0.000)
Execution Overlap	-0.365 (0.025)	-0.447 (0.031)	-0.338 (0.028)	-0.216 (0.736)	0.096 (0.410)
External Financing	0.371 (0.006)	0.286 (0.055)	0.367 (0.006)	1.122 (0.317)	0.148 (0.515)
New Category	-0.063 (0.685)	0.156 (0.464)	0.020 (0.895)	1.596 (0.078)	0.246 (0.231)
Elapsed Time	-0.000 (0.665)	0.000 (0.879)	0.000 (0.961)	0.004 (0.009)	0.001 (0.000)
Funding Period	0.007 (0.101)	0.008 (0.119)	0.007 (0.053)	0.028 (0.172)	0.006 (0.144)
Funding Reward Tiers	0.006 (0.395)	0.007 (0.539)	0.008 (0.296)	-0.023 (0.521)	-0.003 (0.646)
Funding Reward Size	0.000 (0.003)	0.000 (0.000)	0.000 (0.002)	-0.000 (0.779)	0.000 (0.848)
Ln Funding Threshold	0.027 (0.527)	0.048 (0.356)	0.016 (0.693)	0.754 (0.004)	0.157 (0.000)
Ln Funding Exceeded				0.671 (0.008)	0.076 (0.109)
Ln Funding Backers				-0.018 (0.960)	0.074 (0.268)
Baseline Updates				0.172 (0.001)	
Ln Baseline Updates					0.271 (0.000)
Entrepreneur FE	Yes	Yes	Yes	Yes	Yes
Product Subtype FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes
R ²	0.284	0.437	0.279	0.744	0.774
Entrepreneurs	314	314	314	314	314
Observations	722	722	722	722	722

ically, we hypothesize that as entrepreneurs gain experience and encounter increasing unforeseen interdependencies, they miss their predicted delivery date by wider margins. As outlined in Table 2.6, we find that *Project Experience* is positively related to *Delay Indicator* ($p = 0.010$) and *Delay Duration* ($p = 0.001$). As a baseline, with each additional project of experience, the entrepreneur is 11.9% more likely to be delayed. Regarding the magnitude of delay, with each additional project of experience, the average entrepreneur is delayed by an additional 39.6 days. Although *Delay Duration* measures the absolute difference between the entrepreneur's actual and predicted timeline, it is also important to consider the difference on a percentage point basis to account for different predicted project lengths. We also find that *Project Experience* is positively related to *Delay Duration / Predicted Time* ($p = 0.001$). With each additional project of experience the average entrepreneur is delayed by an additional 53.0% relative to her predicted time. Taken together, these findings suggest that, given increasing prediction inaccuracies, entrepreneurs choose to absorb these inaccuracies in the project timeline, leading to increasing delay.

As important context for the above finding, column 4 of Table 2.6 shows that *Project Experience* positively associates with *Predicted Time* ($p = 0.014$). On average, entrepreneurs increase their *Predicted Time* by 8.4 days on each subsequent project. This means that entrepreneurs are not becoming more delayed because they are setting shorter, more aggressive timelines. To the contrary, entrepreneurs give themselves more time on each subsequent project, seemingly anticipating some increase in complexity or adjusting for time they learned that they needed, yet they still miss the prediction by a wider margin. Finally, column 5 of Table 2.6 shows that *Project Experience* is positively related to *Actual Time* ($p \sim 0.001$). On average, entrepreneurs increase their *Actual Time* by 46.4 days on each subsequent project.

To provide an intuitive illustration for interpreting the empirical findings, Figure 2.3 plots the relative trends of *Actual Time* and *Predicted Time* as the entrepreneur gains experience. Figure 2.3 plots coefficient estimates for an alternate non-parametric model of the relationship between experi-

Table 1.4: Delivery and Delay. Ordinary least squares (OLS) estimation. Robust standard errors are clustered at the entrepreneur level. *p*-values are shown in parentheses.

	Delay Indicator	Delay Duration	Delay Duration/ Predicted Time	Predicted Time	Actual Time
Project Experience	0.119 (0.010)	39.616 (0.001)	0.530 (0.001)	8.364 (0.014)	46.386 (0.001)
Failed Campaign Experience	0.139 (0.123)	-2.720 (0.870)	-0.071 (0.749)	15.751 (0.031)	11.345 (0.620)
Prior Campaign Funding Deviation	-0.003 (0.062)	0.094 (0.828)	-0.001 (0.864)	0.095 (0.578)	0.390 (0.411)
Prior Project Delay	-0.042 (0.042)	-21.689 (0.020)	-0.111 (0.690)	4.288 (0.024)	-19.655 (0.034)
Execution Overlap	-0.052 (0.457)	11.587 (0.774)	1.183 (0.233)	-3.993 (0.636)	9.501 (0.831)
External Financing	0.138 (0.196)	95.312 (0.234)	0.371 (0.481)	21.388 (0.376)	124.811 (0.205)
New Category	0.082 (0.611)	39.284 (0.135)	0.425 (0.156)	6.012 (0.546)	46.345 (0.168)
Elapsed Time	-0.000 (0.422)	-0.092 (0.130)	-0.001 (0.353)	-0.023 (0.265)	-0.111 (0.094)
Funding Period	0.009 (0.000)	1.302 (0.059)	0.004 (0.848)	0.147 (0.555)	1.730 (0.022)
Funding Reward Tiers	0.003 (0.454)	0.785 (0.474)	0.001 (0.977)	-0.150 (0.721)	0.996 (0.441)
Funding Reward Size	0.000 (0.807)	-0.011 (0.721)	0.000 (0.707)	0.005 (0.533)	-0.008 (0.830)
Ln Funding Threshold	0.008 (0.786)	22.716 (0.103)	-0.090 (0.533)	19.806 (0.000)	40.817 (0.011)
Ln Funding Exceeded	0.034 (0.195)	28.728 (0.004)	0.201 (0.085)		27.973 (0.014)
Ln Funding Backers	0.008 (0.837)	-34.159 (0.009)	-0.176 (0.281)		-31.278 (0.040)
Entrepreneur FE	Yes	Yes	Yes	Yes	Yes
Product Subtype FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes
R ²	0.662	0.717	0.656	0.752	0.744
Entrepreneurs	306	303	303	314	303
Observations	686	644	644	722	644

ence and the dependent variables by including indicators for project number instead of *Project Experience*.²³ This figure shows that the actual delivery time increases much more sharply relative to the predicted delivery time, with the gap between actual delivery time and predicted delivery time increasing as entrepreneurs gain experience. These empirical patterns are consistent with the theorized project complexity curve and learning curve, respectively, in Figure 3.3.

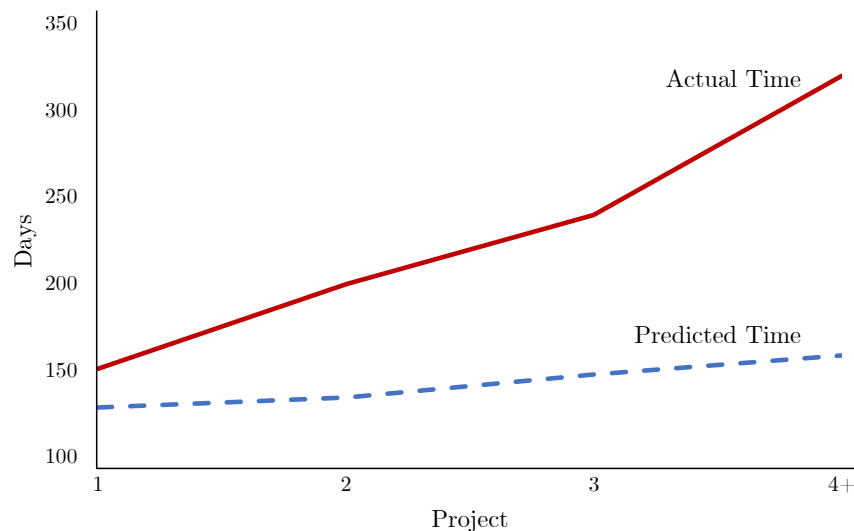


Figure 1.4: Actual Time and Predicted Time with Experience. Visual representation of the relative coefficients of the average actual time entrepreneurs take to deliver a product versus the average time the entrepreneurs predict it will take to deliver, with this relative relationship shown over time as entrepreneurs gain experience. Figure based on coefficient estimates from a non-parametric model detailed in Online Appendix Section A.1.8 that includes indicators for project number instead of *Project Experience*.

1.5.1 LEARNING FROM PRIOR PROJECT DELAY

While *Prior Project Delay* primarily serves as a control variable in the main analyses, we find several statistically significant relationships of note. Recall that *Prior Project Delay* is the duration of the delay divided by predicted time on the prior project, so a value of 1 or 100% means a project was

²³Online Appendix Section A.1.8 details how this figure was created with the additional underlying regressions.

delayed by the same amount of time the entrepreneur predicted the project would take. *Prior Project Delay* is significant and negatively related to *Unforeseen Interdependencies* ($p \sim 0.000$) and its logged value ($p \sim 0.000$). Each increase in days delayed on the prior project equal to the predicted time on the prior project leads to 0.7 fewer unforeseen interdependencies on the current project.²⁴ *Prior Project Delay* is significant and positively related to *Predicted Time* ($p \sim 0.024$). Each increase in days delayed on the prior project equal to the predicted time on the prior project leads to a 4.3 day increase in the predicted time on the current project, showing that entrepreneurs who experience delay give themselves more time on their current project. *Prior Project Delay* is significant and negatively related to *Delay Indicator* ($p \sim 0.042$), *Delay Duration* ($p \sim 0.020$), and *Actual Time* ($p \sim 0.034$). Each increase in days delayed on the prior project equal to the predicted time on the prior project leads to a 4.2% decrease in the chance the entrepreneur is delayed on the current project, a 21.7 day decrease in delay on the current project, and a 19.7 day decrease in the actual delivery time of the current project.

Although we do not formally a priori theorize on the implications of *Prior Project Delay*, the observed patterns have important theoretical implications for this line of research. Based on these empirical findings, we propose two post hoc theoretical explanations. First, it may be the case that learning from a delay, which could be characterized as a failure in a past execution-related prediction, may be more salient to the entrepreneur. From our specific theoretical view, it could be the case that past delays—that come from not accounting for the full set of interdependencies in a prediction—make past unforeseen interdependencies more salient and memorable for the entrepreneur. The mental model of interdependencies she uses to make predictions on the next project would then better account for more of the total interdependencies she would face. Prior work identifies several patterns of entrepreneurs learning from failure of this form (Politis, 2005). Early work by Sitkin (1992) proposes that failure can be especially valuable for learning when the failure is: large enough to draw

²⁴For example, if the predicted time on the prior project is 45 days, then for every additional 45 days of delay on the prior project, the current project would have 0.7 fewer unforeseen interdependencies.

the attention of the entrepreneur, hard to predict, or able to stimulate the entrepreneur to try new ways of doing things. Indeed, a past delay is a notable experience that was hard to predict. On the entrepreneur's next project, this experience can prompt the entrepreneur to account for a previously unforeseen interdependency or at least to give herself more time.

Second, an entrepreneur suffering from a past delay may gain a sense of the shape of the complexity curve that she faces. Based on the premise of our main theory, an entrepreneur in general cannot fully anticipate the full set of interdependencies she will face when she executes her next, more complex project. An underlying assumption here is that the entrepreneur cannot fully anticipate that she cannot fully anticipate the full set of interdependencies. However, through a past delay, perhaps she can learn to anticipate that she cannot fully anticipate the interdependencies. If she does not experience a delay on the prior project, she could still be left assuming that she can anticipate the interdependencies on subsequent projects. To operationalize this awareness that might come from past delay, the entrepreneur could faster-than-linearly build in more extra time into her timeline on subsequent projects. Based on our interviews, most entrepreneurs already try to build in this extra time, but clearly they do so insufficiently. Theoretically, the ideal padding process can be interpreted as an entrepreneur developing an intuitive sense for the shape of the project complexity curve that she faces: she could make predictions of interdependencies based on an extrapolation of the curve, rather than on direct knowledge of the actual interdependencies she will face. This argument mirrors early work by [Toffler \(1985\)](#) and others that theorizes best practices for strategic planning and specifically warns against the pitfalls of straight-line thinking when extrapolating.

1.5.2 SUPPLEMENTAL ANALYSES

We empirically test for and rule out a number of potential alternative explanations that could lead to patterns similar to the main empirical findings or otherwise confound the estimates. First, performance feedback from success or failure on prior funding campaigns could generate an outcome-

aspiration gap for the entrepreneur and affect prediction behavior on subsequent projects. Second, an entrepreneur may base predictions on the relative predictions made by her peers, e.g., predict delivery times that match the average as opposed to predicting how long she actually thinks it will take to deliver. Third, as an entrepreneur gains experience, she may learn that there are limited to no consequences to delaying and, as a result, not care as much about whether she misses the delivery date on subsequent projects. In other words, she may learn that it is “acceptable” to miss delivery dates, especially for more complex projects, which would affect the accuracy of her prediction if it was somehow valuable to promise an aggressive delivery date known *ex ante* to be unrealistic. Fourth, if customers are more likely to fund projects that predict earlier delivery dates, an entrepreneur may be incentivized to overpromise and predict a delivery date that is sooner than her true predicted value. Fifth, higher-quality entrepreneurs may exit the sample when they gain sufficient experience, leaving increasingly lower-quality entrepreneurs in the sample at high levels of project experience. For example, higher-quality entrepreneurs may be able to raise external venture capital financing in lieu of crowdfunding and go to customers through another channel (e.g., direct-to-consumer or retail). Sixth, entrepreneurs who experience a project with significant delay may leave the sample after they “learn their lesson,” resulting in a sample of entrepreneurs who disproportionately do not learn. In this scenario, the failure of missing a delivery time by a large margin could lead better-learning entrepreneurs to pursue an opportunity outside of crowdfunding or to quit altogether. We empirically test each of these alternative explanations and do not find evidence that these mechanisms drive the main results.²⁵

²⁵The Online Appendix provides detailed descriptions of these supplemental analyses with full regression tables. In the order outlined here, the documentation for these analyses appears in Online Appendix Sections A.1.9, A.1.10, A.1.11, A.1.12, A.1.13, and A.1.14, respectively.

1.6 DISCUSSION AND CONCLUSION

This study addresses how an entrepreneur’s experience affects the accuracy of her strategic foresight, reflected in the timeline prediction for a subsequent project. We theorize that as entrepreneurs gain experience, they learn about previously unforeseen interdependencies (which increases the accuracy of subsequent predictions), but they also learn about new opportunities to innovate by implementing new features on a subsequent project, which introduces new, previously unencountered interdependencies (which decreases the accuracy of subsequent predictions). When project complexity increases rapidly, we argue that the latter effect dominates the former, leading increasingly experienced entrepreneurs to make increasingly infeasible predictions. In our crowdfunding context, we show that entrepreneurs take on projects with an increasing number of features and encounter an increasing number of unforeseen interdependencies. In line with our conceptual model, we show that, on average, entrepreneurs miss their predicted timeline by a gap that grows by nearly six additional weeks (an additional 53.0% relative to their predicted timeline) on each subsequent project.²⁶

1.6.1 STRATEGIC FORESIGHT UNDER COMPLEXITY

By taking the view that accurate strategic foresight depends on anticipating complexity in a strategy, we put forth an alternative explanation for the widespread challenge of making accurate predictions in entrepreneurial settings. In a review, [Townsend et al. \(2018a\)](#) note that there has been sparse work on understanding entrepreneurship in a complex environment where the construct of “uncertainty

²⁶Our findings are reminiscent of the conventional managerial wisdom embodied by the Peter Principle ([Peter & Hull, 1969](#)), often phrased as “Employees rise to their ‘level of incompetence’ in a hierarchy,” i.e., managers who are promoted due to success in a prior job are then confronted with managing a new set of responsibilities unrelated to what made them successful previously. One could summarize the findings of this study as, “Entrepreneurs rise to their ‘level of incompetence’ in strategic foresight,” i.e., entrepreneurs who succeed in making a product and then continue to add features and increase the complexity of that product, are required to manage systems of interdependencies which they have never encountered previously and which they are ill-equipped to manage.

has been stretched to try to address aspects of unknowingness that are better conceptualized as complexity” (p. 674). We seek to address this gap by taking a view that complexity presents a barrier to what can and cannot be learned, acting as an important constraint on the returns to experience. We assert that it is the challenge of accounting for complexity, rather than just uncertainty from a lack of available knowledge, that limits how much an entrepreneur can improve her ability to make predictions. Two important implications emerge when accounting for the role of complexity in strategic foresight.

First, the nuanced complexity-based mechanism we propose stands in contrast to the view of the extant literature that entrepreneurs’ prediction inaccuracies stem from a general characterization of entrepreneurs as being “overconfident.” Prior literature documents compelling evidence that entrepreneurial entrants make infeasible predictions about their own future performance, leading them to enter markets they should not (e.g., [Artinger & Powell, 2016](#), [Cassar, 2010](#), [Chen et al., 2019](#), [Forbes, 2005](#), [Hayward et al., 2006](#), [Wu & Knott, 2006](#)). This literature generally frames this observed pattern as a consequence of entrepreneurial overconfidence ([Camerer & Lovo, 1999](#)). However, we assert that the inherent complexity involved in entrepreneurial strategy and new product development may be a key antecedent to what otherwise appears as overconfidence. Our view aligns with [Hogarth & Karelaia \(2012\)](#), whose simulation model shows how over-entry can occur among both overconfident and underconfident entrepreneurs. While an entrepreneur’s lack of a full understanding of the complexity she faces may appear as overconfidence to an observer, the natural emergence of complexity likely accounts for at least some of the error in her predictions.

Our study documents empirical evidence for this nuanced characterization of entrepreneurs facing complexity rather than being generically overconfident. The entrepreneurs in our setting accumulate information through experience that should help address an overconfidence bias that stems from a lack of information. However, inconsistent with a basic overconfidence explanation, we find that entrepreneurs actually become less accurate as they accumulate knowledge from experience. That

said, our arguments do not rule out the possibility that overconfidence still exists.

Second, accounting for the role of complexity implies that learning-from-experience may not be a cure-all solution for inaccurate strategic foresight. In a seminal study of the automobile industry, [Levitt et al. \(2013\)](#) show that model changeovers disrupt the learning curve; when firms add new model variants, prior learning is less helpful. So the outstanding question is why that is the case? We propose that it is the emergence of new complexity unrelated to prior experience that impairs strategic foresight.

As entrepreneurs learn from prior projects, an important manifestation of this learning is to add new features to their products, which in turn drives complexity that impairs strategic foresight. [Ethiraj et al. \(2012\)](#) find that firms face pressure to address customer requests with incremental product innovations, but even incremental changes can precipitate a cascade of impacts across interdependent parts of the product and organization. The entrepreneurs we study face this exact challenge, with severe consequences for the accuracy of their strategy foresight. Under the assumptions that entrepreneurs inevitably face this incentive to improve over time and that complexity is difficult to address and anticipate, the unfortunate implication is that strategic foresight will face a perpetual headwind.

1.6.2 STRATEGIC FORESIGHT AS MULTIPLE PREDICTIONS

To empirically study strategic foresight, we make a key advance with our explicit interpretation of strategic foresight as not just one prediction but the combination of a set of predictions. In contrast, the broader set of work on strategic decisions—and specific studies on foresight—focus on whether a manager or another actor can make a sole prediction or decision, e.g., enter into a market ([Camerer & Lovo, 1999](#)) or invest in a specific firm ([Csaszar & Laureiro-Martínez, 2018](#)). We argue that strategic planning must inherently invoke several predictions simultaneously, whether articulated or not, because the predictions depend on one another. Predicting a value proposition also entails predicting a cost structure for delivering that value proposition such that the aggregate strategy is viable.

At a general level, our findings suggest that as entrepreneurs gain experience across projects, their strategic predictions on a successive project will be less accurate; but that characterization would be far from a complete story. We show a trade-off among the aggregate predictions that comprise strategic foresight broadly (Ethiraj & Levinthal, 2009, Talbot, 1982). We gain several advantages by studying product specification, delivery timeline, and complexity simultaneously. In our study, entrepreneurs pursue success in achieving the predicted product specification but at the cost of delivering their product on time. But that is a choice they made. In principle, the entrepreneur could deliver the product on time but at a lower value proposition than they originally predicted.

There are many high-profile examples of entrepreneurs prioritizing the delivery of an initially specified product over staying within the initially predicted timeline. In July 2017, Elon Musk promised Tesla would deliver 20,000 Model 3 cars in December of that same year. However, Tesla only produced 2,425 cars the entire fourth quarter of 2017, falling short of Musk's prediction by 93%. Tesla eventually reached their predicted product specification, and even exceeded it, reaching over 10,000 vehicles per week in 2018, but far behind the initially predicted timeline. Speaking of Tesla's tradeoff between delivering on predicted product specification versus timeline, Musk himself said, "It pretty much always happens, but not exactly on the time frame."

By considering several predictions simultaneously in a holistic view of strategic foresight, future research can provide more nuance in the ways in which entrepreneurs who have previously been categorized one-dimensionally as failures (or successes) might actually have succeeded (or failed) along other overlooked dimensions. In doing so, we might show that some failures are driven by an intentional choice to succeed on other dimensions. By recognizing the multiple predictions inherent in strategic foresight and their relative prioritization, entrepreneurs and investors may be able to improve performance. For example, given that additional costs may be required to achieve a fixed product specification, especially as entrepreneurs gain experience from past projects and take on increasingly complex projects, future research could explore how entrepreneurs and investors can identify situa-

tions where adjusting the product specification may be preferable to accruing high costs or missing delivery dates.

1.6.3 PREDICTING TIMELINES: STRATEGIC IMPLICATIONS

This research brings attention to the strategic problem of predicting and managing timelines. In the context of product crowdfunding, we show that entrepreneurs—even and especially those with experience—consistently struggle to predict accurate timelines. This struggle extends well beyond our context to entrepreneurs generally. For instance, consider Chinese electric vehicle company Faraday Future, which initially predicted it would begin production on its flagship SUV in 2018. When this timeline turned out to be wrong, the company raised an extra \$225M in emergency bridge financing in order to keep the company alive and get the company and its product to where they needed to be for a future public offering. While Faraday Future recovered from its poor timeline prediction, many other electric vehicle startups were not so fortunate: early pioneer Fisker Automotive was forced to shut down due to a poor timeline prediction that led to them running out of cash before being able to raise more money.

Large, established firms face this same timeline challenge. Apple missed its predicted timeline to ship the HomePod in 2017, AirPods in 2016, and the Apple Watch in 2015. Similarly, Microsoft missed its predicted timeline to release many of its new operating systems, to ship Surface Earbuds in 2019, and to push a security update in 2017. Missing predicted timelines is also common in other settings such as big box office movie releases. In two high-profile examples, unanticipated post-production complexity led to the delayed release of *Titanic* (from July to December of 1997) and *Gravity* (from November 2012 to November 2013).

Based on our theoretical framework and our empirical observations, we propose three ways a manager can make more accurate timeline predictions for firm strategy. First, we see an opportunity for firms to make an intentional effort to better anticipate the non-linear nature of complexity by account-

ing for the faster-than-linear increase in unforeseen interdependencies when building out projects. In our setting, both unforeseen interdependencies and delays are ubiquitous and increasing over time. Just as becoming aware of personal biases or tendencies towards overconfidence can allow managers to make better decisions (Pope et al., 2018, Lee & Huang, 2018), becoming self-aware of the true realities of complexity could theoretically empower managers to make more accurate timeline predictions.

To put this argument in more colloquial terms: we all face known knowns, known unknowns, and unknown unknowns. The interdependencies that an entrepreneur could face fall into these buckets. Through experience, it is theoretically possible that an entrepreneur could become more aware of the rate at which unknown unknowns will arise—in essence, allowing an entrepreneur to treat them as known unknowns—and to account for those unknowns when making predictions through extrapolation based on a higher level of intuition for the shape of the project complexity curve. How does one take this to practice? As one suggestion, we document that entrepreneurs already engage in a process of padding their timelines with extra time, albeit to an insufficient degree. We recommend that entrepreneurs act on the insight of this research by padding timelines with the complexity curve in mind, way more than their prior (incorrect) intuition would suggest.

Second, firms can learn to make more accurate timeline predictions by internalizing salient experiences with interdependencies. We find that the delay on a subsequent project is partially offset by experiencing delays on the prior project. Thus, previously delayed firms have a unique opportunity to carefully identify the specific unforeseen challenges that contributed to the delay and then to account for those factors when making subsequent predictions. Of course, we would not suggest that firms intentionally cause a delay in pursuit of this benefit. But certain circumstances allow firms to engineer controlled experimental experiences that make unforeseen interdependencies salient like a prior project delay but without the same costs. For example, we spoke with an Apple manufacturing manager about how they now stress-test prototypes and test-run manufacturing small batches to try to identify interdependencies before a high-stakes product launch.

Third, future research should explore whether firms can improve timeline predictions in complex situations by increasing their knowledge diversity (Olson et al., 1995, Keller, 2001). A more diverse knowledge base increases the breadth of interdependencies a firm will be aware of when making predictions. An increased breadth of awareness should lead to improved foresight (Csaszar & Laureiro-Martínez, 2018). As the fundamental challenge of complexity, a firm cannot ex ante anticipate all the relevant interdependencies—which means the firm cannot ex ante plan for which knowledge it will need. Thus, there may be value in intentionally maintaining a diverse set of experience at the table beyond what the firm ex ante expects to be directly relevant to a given project: increasing diversity could increase the chances that someone will anticipate a relevant interdependency. If firms only hire or seek input from a narrow set of people assumed to be relevant, the value of the marginal voice for identifying unknown interdependencies is limited. Diversity could also be sought outside the boundaries of the firm (Aggarwal et al., 2020). For example, crowdsourcing (e.g., open innovation tournaments) provides access to more diverse knowledge and improves performance when searching for the global optimum (Afuah & Tucci, 2012). While the current study does not empirically measure this channel, future work could directly measure the impact of knowledge diversity on timeline predictions.

2

Complexity and Timing in Novel Ventures

2.1 INTRODUCTION

Strategy and entrepreneurship research traditionally highlights the positive returns to pursuing a novel, or new, opportunity. Under this view, novelty in resources employed (Felin et al., 2023, Ahuja & Katila, 2004, Barney, 1991), activity configuration (Ott & Eisenhardt, 2020, Porter & Siggelkow, 2008), and market positioning (Kim & Mauborgne, 2014) contribute to competitive advantage. Schum-

peter (1939) proposed the now-canonical notion of knowledge recombination—the combination of existing knowledge into novel combinations—as a critical process by which entrepreneurs create or discover novel opportunities with the potential for value. Since then, an extensive body of empirical work has identified the circumstances fostering recombination and when this process results in novel products, services, and business models (e.g., Gruber et al., 2012, Keijl et al., 2012, Carnabuci & Operti, 2013, Mukherjee et al., 2016). Consider Spotify, the world’s most popular internet music subscription service in 2023. At its founding in 2006, it represented a novel combination of existing ideas: providing music over the internet, previously pioneered by Napster¹, Apple iTunes, and others; a freemium model offering customers a choice between advertising and subscription, previously pioneered by video game companies like Epic Games; and social interaction, previously pioneered by social networks like MySpace (Tidhar & Eisenhardt, 2020).

However, the literature on recombination largely overlooks the process by which entrepreneurs translate a potential opportunity in the abstract sense into a tangible business in the concrete sense. In short, this literature on knowledge recombination overlooks its implications for execution and how long it takes to execute on a novel recombination. While scholarship documents and celebrates how entrepreneurs identify abstract opportunities (Alvarez & Barney, 2007), it does not yet address the execution needed to bring that idea to market, which represents, in our view, the most substantive role of the entrepreneur and takes real time and perseverance to see through. More recent work alludes to the importance of execution—Adner & Levinthal (2008) suggest that all organizational activities (including exploration) are fundamentally exploitative, and McDonald & Eisenhardt (2019) highlight many different actions that organizations must take to actually create value—but does not address entrepreneurial execution directly. Understanding the relationship between the pursuit of a novel recombination and its execution is important because it allows us to understand why otherwise ap-

¹Spotify co-founder Daniel Ek was previously the CEO of uTorrent, which made money by monetizing pirated music and movies.

pealing new business opportunities fail to materialize in the marketplace. Looking back at the history of Spotify, it is striking both how long it took to set up the imagined service and how the novelty of the business itself is what led to the executional challenges that prevented it from launching sooner. It took years to secure global licenses to stream music from copyright holders, a challenge exacerbated by the need to track royalty payments for streams by users at different tiers of the freemium model and the need for technology to protect the IP of copyright holders when users engage socially (Cozzolino & Rothaermel, 2018). Fortunately for music listeners, Spotify eventually persevered, but as we will document and argue in this study, many firms pursuing novel recombination do not, and where they do, it can take a long time, longer than anticipated, for these novel opportunities to realize value as businesses.

Taking time and execution into account, we theorize that the pursuit of novelty, despite the strategic benefits described in the literature, comes with significant downsides that call into question whether its pursuit is justified. To conceptualize the relationship between a novel opportunity, formed through recombination, and the time it takes to execute on that opportunity and translate it into a tangible business, we take a complexity-based view of entrepreneurship. We take the view that entrepreneurial execution is fundamentally about managing complexity, i.e., interdependencies among components or people (Brown & Eisenhardt, 1997, Simon, 1962). This perspective generates a number of implications for entrepreneurial strategy and, specifically, the time needed for execution and the challenges of predicting that timeline. Complexity non-linearly increases the time needed for execution (Sommer & Loch, 2004, Anderson, 1999), and its unforeseeable nature makes it more likely that entrepreneurs (Peterson & Wu, 2021) and venture capital firms (Huang, 2017, Csaszar, 2018) underestimate the time needed for execution. We argue that novel recombination increases the need to manage complexity because interdependencies not previously anticipated or dealt with exacerbate the need for more time and the challenges of predicting that time. This makes it more likely that an entrepreneur will face an unanticipated need for more time to continue executing beyond what exist-

ing financing can support, elevating the risk that the firm will run out of money. Without emergency financing, it would then have to shut down even if the opportunity pursued is otherwise promising.

In a study of 31,450 entrepreneurial firms across 94,549 venture funding rounds, we find compelling empirical support for the relationship between novel recombination and complexity, timelines, and performance. We first explore the relationship between novel recombination and complexity and find that doubling novel recombination increases the number of references to managing complexity in employees' job descriptions by 18.7% relative to the sample mean. We then show how executing on any novel recombination requires more time for firms to grow in value: \sim 1 month longer to reach a \$100M USD valuation and \sim 6 months longer to reach a \$1B USD valuation. Consistent with our theory of novel recombination leading to unanticipated resource needs, we also find that firms with novel recombination are more likely to require bridge financing or shut down. Doubling novel recombination increases the likelihood that the firm will need bridge financing by 2.5% and increases the likelihood a firm will shut down by 3.5% (both relative to the sample mean).

This study makes three important contributions. First, we put forth a complexity-based view of entrepreneurship that highlights resolving interdependencies as the primary role of the entrepreneur. This is distinct from an uncertainty-based view (Knight, 1921) and has implications for entrepreneurial experimentation (Camuffo et al., 2020). Second, we highlight the importance of considering firm-level timing and the consequences of delayed value realization. This stands in contrast to existing theoretical work on S-curves (e.g., Suarez & Lanzolla, 2007) and empirical work on patents (e.g., Fleming & Sorenson, 2001). We take a firm-level perspective that accounts for the gap between an idea (i.e. patent) and value realization via entrepreneurial execution. Third, we emphasize the importance of understanding the determinants of entrepreneurial failure and bring attention to the ubiquitous yet heretofore under-studied phenomenon of bridge financing. While much of the existing literature focuses on the right tail of entrepreneurial success, as Artinger & Powell (2016) and others illustrate, there are key lessons to be gleaned by studying entrepreneurial failure. That said, while the focus of

this paper is entrepreneurial failure, some of our results also suggest that novel recombination does lead to an increased likelihood of IPO and higher valuations in the long run. However, our point is that the road to this success runs through the graveyard of firms that did not make it to the end.

2.2 THEORETICAL BACKGROUND

We outline a conceptual perspective that highlights the implications of novel recombination by distinguishing the notions of entrepreneurial opportunity and entrepreneurial execution. Whereas recombination generates entrepreneurial opportunities in the abstract, realizing the potential opportunity requires execution, a concrete process that takes place in the real world involving the actual work of individuals and organizations.

At the start of this section, we briefly outline the existing perspective of opportunities as emerging from the novel recombination of existing knowledge. We emphasize that this process is abstract and largely conceptual, in contrast to the process of execution.

After briefly summarizing this view on opportunity identification, this section then focuses in depth on our view of entrepreneurial execution as managing complexity, which we refer to as a *complexity-based view* of entrepreneurship. We argue that novel recombination in the abstract requires entrepreneurial execution to realize the potential opportunity in concrete practice, and that this execution requires *managing complexity*. We argue that managing more complexity takes more time and actually drives time-compression diseconomies. As such, increasing complexity increases the time required to execute on the opportunity.

In the subsequent *Hypothesis Development* section, we revisit the overarching link between novel recombination and firm performance. Thus, we outline an important conceptual relationship between recombination, the time required to actualize an abstract opportunity, and the impact of this time requirement on firm performance.

2.2.1 OPPORTUNITIES AND RECOMBINATION

Existing literature establishes that entrepreneurs recombine knowledge in the pursuit of creating or discovering potential opportunities. Schumpeter (1939) first established this now-canonical notion of recombination as the critical process by which entrepreneurs create or discover novel opportunities. More recent work emphasizes the importance of newness or novelty in recombination. Fleming (2001) defines a recombination as “either a new combination of components or a new relationship between previously combined components” (p. 118). Alvarez & Barney (2007) emphasize that entrepreneurs identify a “new opportunity” comprised of “new products or services,” with the qualification of being *new* appearing on every page of the article. While this literature recognizes two distinct views of opportunity identification (endogenous vs. exogenous), both views clearly concur that novelty is the substance of the opportunity, and “the creation of any sort of novelty... consists to a substantial extent of a recombination” (Nelson & Winters, 1982).

In other words, combining two heretofore disparate concepts creates something novel, which could afford entrepreneurial opportunity. For example, although electric motors and mechanized vehicles (e.g., steam) each have a long independent history, the idea of recombining the concept of an electric motor with the concept of mechanized vehicles leads to the novel concept of an electric vehicle, with the first patents covering this novel recombination appearing in the late 1800s. Of course, this conceptual exercise merely identifies, not delivers, the opportunity of an electric vehicle.

In particular, entrepreneurs often seek out *novel* recombinations due to their potential to create value (e.g., Amit & Zott, 2001, Singh & Fleming, 2009, Kneeland et al., 2020). The extant literature establishes a connection between novelty and the existence of valuable opportunities (Zott & Amit, 2007, 2008). This literature follows naturally from the notion of Schumpeterian rents, where novel recombination is the engine of competitive imperfection. Here, novelty can disturb the general economic equilibrium and thereby provide the opportunity for supernormal profits, which persist until

imitation occurs thereby diminishing and ultimately eliminating the novelty as equilibrium is restored (Schumpeter, 1939).

We will revisit novel recombination in the *Hypotheses Development* section and link it to complexity and timing. For now, we focus on these concepts of complexity and timing, as they require detailed explanation before we can turn back to novel recombination later.

2.2.2 EXECUTION AND COMPLEXITY

While we take novel recombination as a conceptual process to identify entrepreneurial opportunities, all opportunities must, of course, be executed on to realize their potential. Novel recombination in opportunity identification is *ex ante* to the material realization of potential value. Execution is required to actually deliver the imagined product or service and to actually collect any rents. In other words, opportunities only have potential energy, whereas execution makes that potential kinetic.

We take the view that entrepreneurial execution is fundamentally about managing complexity. Entrepreneurial execution and complexity are so intertwined that it is difficult to construct an example of an entrepreneurial endeavor that does not require managing complexity. Indeed, we view managing interdependencies as the very value-add of the entrepreneur. For example, consider an entrepreneur executing to build a payments platform, to develop a new drug, or to bake the most delicious cookies. In each case, the critical challenge of execution is managing the interdependencies between the customer and various banks, between the drug and the patient, and between all the various ingredients, respectively. Even the most basic day-to-day tasks of an entrepreneur revolve around managing complexity. When an entrepreneur sends an email, picks up the phone, unpacks a delivery and moves it to another box, or performs many other routine activities, the purpose is to manage an interdependency. If an entrepreneur employs an engineer and a salesperson, the very purpose of the entrepreneur is to manage the interdependencies between the two roles. It is hard to even articulate an example of an entrepreneurial endeavor that does not involve managing complexity, but instead only consists of a

singular and totally modular component in isolation. In that scenario, the entrepreneur has no purpose nor value-add. Therefore, we take a complexity-based view of entrepreneurship, namely that the fundamental act of entrepreneurship is execution, which requires managing complexity.

Importantly, viewing entrepreneurship through the lens of managing complexity is distinct from the ubiquitous view of entrepreneurship as managing uncertainty (Knight, 1921, Coase, 1937, Keynes, 1937). As explained in Townsend et al. (2018a), problems of uncertainty relate to an unknown relationship between only two variables in an isolated system (e.g., the relationship between a single switch and a single bulb), whereas problems of complexity are defined and driven by both the multiplicity of variables and the interdependencies between those variables (Simon, 1962, Porter & Siggelkow, 2008) (e.g., the relationship between many interdependent switches and a series of bulbs). The *Discussion* section further explores this complexity-based view of entrepreneurship and its implications in practice.

2.2.3 STRATEGIC IMPLICATIONS FOR TIMING

We argue that managing complexity during execution is challenging and takes meaningful time to achieve. If we take an NK landscape perspective, as complexity increases, the ruggedness of the landscape increases, and the search time required to find the optimum increases non-linearly. As a result, in complex systems the relationship between time required and opportunity realization is not constant but, rather, non-linear (Anderson, 1999).

The concept of time-compression diseconomies is intrinsically linked to the non-linear escalation of time demands imposed by complexity (Dierickx & Cool, 1989, Scherer, 1967, Srikanth et al., 2021). This concept posits that as complexity in developing resources intensifies, the efficiency losses associated with attempts to shorten the time frame for development become increasingly pronounced. The root of these inefficiencies lies in the inherent characteristic of complex systems—the irreducibility of interdependencies (Sommer & Loch, 2004). Modular components can be reduced and separated, but

interdependent components are irreducible (Baldwin & Clark, 2000). Given irreducible complexity, there is some fundamental, minimum time required for execution that cannot be bypassed without risking the integrity of the pursued opportunity. This minimum time, expanding with the degree of complexity, underlines the limitations to substituting temporal resources with monetary investments (Talbot, 1982). In other words, when managing complexity entrepreneurs cannot just buy lightning-fast execution.

Building on the understanding of the irreducibility of execution time when managing complexity, it becomes evident that escalating levels of complexity invariably lead to delays in realizing the entrepreneurial opportunity. This relationship represents another argument for the existing theory of complexity as a source of competitive advantage (Rivkin, 2000). More complexity widens the temporal gap between innovation and its potential imitation by competitors. This delay in imitation is instrumental: it serves as a strategic buffer, granting an original innovator a temporal window of advantage. Complexity's impact on timing is at the root of both its role in challenging execution efficiency and its role in bolstering competitive positioning.

We argue that the execution time required to manage under complexity represents a challenge throughout the firm life cycle. First, at inception, an entrepreneur must set up a system that correctly captures relevant interdependencies to even offer the basic product or service at all. This concept builds on the concept of an "incubation period" put forth by Moeen & Agarwal (2017). But even after having a working product or service, an ongoing need to manage more interdependencies persists, or even intensifies, as the firm scales. From inception to scale, managing more complexity takes more time (Peterson & Wu, 2021).

PRODUCT DEVELOPMENT: SETTING UP A COMPLEX SYSTEM

Complexity makes it take longer to develop the product that delivers on an opportunity's potential value: development is incomplete until after addressing a preponderance of, or even all, necessary in-

terdependencies. Sommer & Loch (2004) theorize that there are irreducible interdependencies that must be resolved to offer a product or service which cannot be foreseen. Given that execution time increases with complexity, and that complexity is unforeseen, an entrepreneur faces the risk of underestimating the needed time and ends up delayed.

Mistaken timelines are problematic for entrepreneurs and investors because value emerges late in complex systems. When building a mousetrap piece by piece, the trap does not work and performs no meaningful function until the final piece is correctly put in place to resolve the final interdependency, at which point the trap suddenly goes from completely nonfunctional to fully functional.² Entrepreneurial execution involves transferring an abstract blueprint of opportunity to a physical product or acting organization that can actually deliver the promised potential: connecting the final wire in the circuit so the light turns on, obtaining the final required legal licence to operate, fixing the final bug in the code, hiring the key role, etc. Any realization of the opportunity's potential value is delayed until the final interdependency in the minimum-reducible product system is resolved. Even in the most basic and quintessential juvenile entrepreneurial endeavour of a lemonade stand, value cannot be realized until the full set of interdependencies is resolved: between the location, traffic, lemons, water, weather, etc. Even more so than the lemonade stand, take the example of FedEx. FedEx sought to execute on the opportunity to ship something from anywhere to anywhere. This required resolving a massive, complex and costly system of interdependencies. Here, the irreducible system to realize the opportunity's potential value required building an entire fleet of coordinated planes to take off and land in order and on schedule. This execution required a lot of time, much more than anticipated, and proved so costly the company faced the risk of bankruptcy multiple times before achieving the

²This is reminiscent of the notion of a tipping point, where reaching a certain critical mass or point in the process suddenly leads to significant and meaningful change (Navis & Glynn, 2010). Tipping points also play out in network effects when there are increasing returns to scale. For example, in the war between HD-DVD and Blu-ray, the two players battled for years until Blu-ray crossed a critical threshold at which point the value of Blu-ray shot up and the value of HD-DVD went to zero. Importantly, increasingly complex systems have an increasingly narrow set of configurations that will trigger the tipping point.

vision of a global overnight delivery system.

This complexity-based view of entrepreneurial execution stands in contrast to an uncertainty view of entrepreneurship, where resource-light experimentation can quickly provide relevant information useful for the realization of the opportunity. Cerebral exercises of information gathering do not help manage the challenge of complexity. Using the vernacular of the NK literature (Kauffman, 1995, Levinthal, 1997), when levels of interdependence are high and the entrepreneur finds herself on a rugged landscape, step-wise experimentation is of limited value in searching for the peak opportunity. Phrased differently, in complex systems the cost of experimentation can be preclusively high.

As detailed in the *Discussion* section, the reality of high search cost on complex rugged landscapes is why the Silicon Valley mantra to “fail fast” through rapid experimentation is often inconsistent with the resolution of complexity. The long execution time required to manage complexity means that giving up “fast” may not be possible without also giving up before reaching a point where the opportunity can be realized. For example, what would it have meant for FedEx to fail fast? If they had run an experiment with a single plane and failed to capture any value, could they have then concluded the opportunity was not valid and given up? No, because the whole business does not work until they have worked through the complexity of mobilizing a fleet of planes at scale that are coordinated with one another. FedEx could only resolve the problem of complexity through execution, which required building the fleet, solving the coordination challenges of takeoffs and landings, etc. Although scaling up by adding new routes certainly increased the complexity of FedEx over time, there was some baseline amount of complexity inherent in the initial minimum network of planes required to perform the single offering of overnight shipping from anywhere to anywhere that the firm had to expend time managing prior to adding any scale.³

³Or consider the example of speech recognition, where the value of this technology increased slowly despite massive investments of time over decades of product development. It was not until the complex problem of speech recognition was resolved to the point of near-human accuracy that the value could be realized. Only after the large amount of time invested allowed the technology to clear the threshold of human-level accuracy did the applications and value increase very rapidly.

SCALE: GROWING AND MAINTAINING A COMPLEX SYSTEM

Initially resolving the irreducible set of interdependencies to offer the basic value proposition of an opportunity generally only serves a small subset of potential customers. However, producing a single prototype is very different from producing thousands of units, which is very different from producing millions of units. As the number of units and/or customers increases, firms need to codify processes, source sufficient inputs, manage more relationships, separate out parts of the process across teams, etc. Operating at scale requires addressing increased complexity in the organizational structure through increased coordination between parties, straining managers tasked with managing those interdependencies (Penrose, 1959, Cyert & March, 1963, Chandler, 1966). That said, the extant literature also suggests that scale can also act as a barrier to entry and a source of competitive advantage (Yao, 1988, Lieberman, 1989). But our view is that scale provides advantage because of the inherent time-compression diseconomies generated by complexity. If scale were simply a matter of doing the same thing again and again, scale would be readily imitable. Thus, the challenge and advantage of scale derive from the time required to address interdependencies introduced as the number of units produced and/or customers served increases.

We assert that the greater complexity from building scale increases execution time and can lead to increasing delays in execution if the greater complexity is unforeseen. For example, as Tesla ramped up production of its Model 3 car and increased scale, it repeatedly missed production targets as it struggled to resolve interdependencies with the outside battery suppliers and subcontractors to provide sufficient supplies for a larger number of units. The rollout of the Pfizer-BioNTech COVID-19 vaccine experienced delays in some regions due to complexities in manufacturing, distribution, and meeting the various regulatory requirements across different countries. Or consider Qualcomm's 3G technology. Even though the technology was fully developed, the value of Qualcomm's investment in 3G could not be realized at scale until after the interdependencies with the telecom operators, hand-

set manufacturers, and end users were resolved. Managing this complexity took time, and the value of the opportunity was delayed until all the relevant interdependencies could be resolved. Clearly, the interdependencies involved in building scale take time to manage.

2.3 HYPOTHESIS DEVELOPMENT

We now turn back to the notion of novel recombination, which will be the main independent variable of interest. Prior to this point, our theoretical discussion has largely centered around building up our perspective on the relationship between execution, complexity, and timing. Using this theoretical perspective, we can now make predictions about the implications of novel recombination on firm performance.

2.3.1 NOVEL RECOMBINATION AND COMPLEXITY

We first establish the relationship between novel recombination and complexity. Novel recombination is a theoretical exercise in which two heretofore disparate concepts are brought together. Realizing this conceptual union in the real world requires execution to manage real interdependencies. For example, the novel recombination between *electric* and *car* creates a host of novel interdependencies. The product contains new interdependencies between the electric battery and the motors to turn the wheels, between the charging port and the battery, etc. The organization must also manage new interdependencies between the company and the lithium, nickel, and cobalt suppliers required to make the batteries, between the company and third-party charging infrastructure, etc. Novel recombination leads to novel interdependencies, which are the building blocks of complexity. As a result, novel recombination leads to novel complexity.⁴

⁴In some ways, novel recombination and novel complexity are related by construction. In a novel recombination, two things that have never been combined previously are put together and connected for the first time. This new dyad is, by definition, a novel interdependency at an abstract level, which then translates to complexity at a concrete level.

One key place where complexity emerges as a result of the pursuit of novel recombination is in the firm's activities (Porter & Siggelkow, 2008). Increasing novel recombination leads to increasing novel interdependencies, which are encapsulated by specific firm activities managed by specific roles within the organization; as novel recombination increases, an increasing number of a firm's activities will involve managing novel interdependencies. Take the example of Carta, a cap table management software company founded in 2012 to manage company ownership and issue digital stock certificates to investors (pursuing novel recombinations relative to company ownership, legal compliance, the cloud, etc). In practice and in order to realize the opportunity inherent in its novel recombination, Carta built out a system of activities to execute and manage the relevant interdependencies. This system required hiring specific employee job roles to perform specific activities to manage specific interdependencies. For example, the "Engineering Manager of Investor Services" manages the interdependencies between the software product and the venture capital operations teams. Carta organized software roles to handle relevant components of its novel recombination with coders handling the interdependencies between contracts and software, valuations and software, investor services and implementations, etc. Carta's specialized and diverse set of software roles and related activities stands in contrast to the more streamlined set of roles and activities that would be required for a firm with less novel recombination.

Novel recombinations require managing novel interdependencies, which in turn increase the complexity of the firm's organizational activities. Phrased differently, realizing a novel recombination requires executing on and managing more complexity.

Hypothesis 3 *Pursuing more novel recombination leads to an increasing need to manage complexity.*

2.3.2 NOVEL RECOMBINATION AND TIMING

This recombination-complexity relationship has meaningful implications for the time it takes to realize an entrepreneurial opportunity. We highlight the important but overlooked link between novel

recombination and the duration of execution needed to realize the opportunity in practice. While existing literature predominantly addresses how knowledge recombination contributes to the genesis of abstract ideas, it often presupposes immediate realization of the opportunity's value. For example, the literature on the value of novel recombination measured through patents (Fleming & Sorenson, 2001, Fleming, 2001, Aggarwal et al., 2020) assumes that the patent in and of itself is the source of value; once the patent exists, the value also exists. Or work that measures novel recombination in research papers (Uzzi et al., 2013) assumes that the value of the recombination is inherent in the abstract idea. That assumption might apply in the context of academic research, but we argue much less so in entrepreneurship. This literature overlooks that there is a gap between the abstract idea and realizing its value in practice. As we will argue, novel recombination increases the gap between the abstract idea and practical application, with implications for firm performance.

Novel recombination drives novel, and thus unforeseen, complexity, which, due to the irreducibility of complexity and time-compression diseconomies, delays the realization of an opportunity. Figure 3.3 outlines the amount of time it takes to realize value at varying degrees of novel recombination. Here, value corresponds to the rents or supernormal profits the entrepreneur is able to extract from a given opportunity. As the amount of novelty and the associated complexity in execution increase, the total amount of time required both to develop a product and to scale increases. If value cannot be realized until the critical, irreducible set of interdependencies is resolved, the realization of value will be delayed in complex systems.⁵ In the extreme, this would look like a step function where no value is assigned to an opportunity for an extended period of time until all the interdependencies are resolved,

⁵For example, consider Firm A which must resolve five interdependencies and Firm B, pursuing more novel recombination, which must resolve ten interdependencies. If both firms resolve five interdependencies, the value of Firm A will go up because all five interdependencies are resolved; however, the value of Firm B will not increase because half of its "mousetrap" is still not assembled, so no value can be realized given these outstanding interdependencies. In other words, the more novel recombination and inherent complexity pursued by the project, the smaller the slope of value over time will be early on. As systems become more interdependent and less modular, value will increase in an increasingly step-wise manner and, in the extreme case when the system is fully interdependent, the slope will be flat with no increase in value until all the interdependencies are resolved (at which point there would be a step-wise increase in value with an infinite-slope vertical line jump).

after which the full value of the recombination emerges.

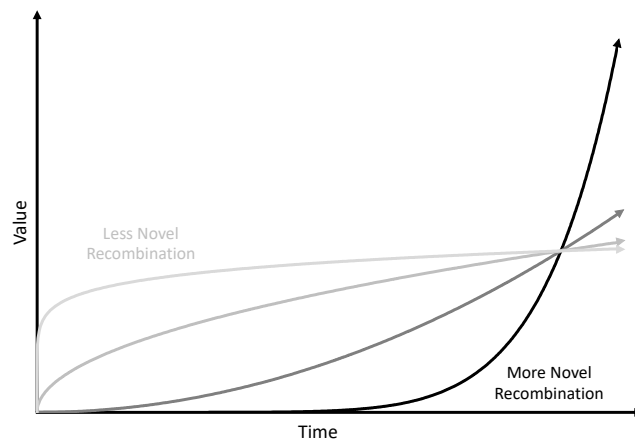


Figure 2.1: Novel Recombination and the Timing of Value Realization. The vertical axis *Value* represents the total value of the opportunity. The horizontal axis *Time* represents the total time allocated to executing and resolving the complexity and associated interdependencies.

Hypothesis 4 *Pursuing more novel recombination requires increasing execution time to realize the potential value.*

2.3.3 NOVEL RECOMBINATION, DELAY, AND FAILURE

The increasing time demands of complexity and the associated delays in realizing the opportunity in practice can pose an obstacle for entrepreneurs and for venture capital investors. Given three key assumptions, this complexity in novel recombination can lead firms to experience unforeseen delays and often fail. First, as we discussed previously, complexity is irreducible and requires significant execution time to manage. Second, complexity is also unforeseeable, particularly when there are novel interdependencies driven by novel recombination (Kapoor & Wilde, 2023, Peterson & Wu, 2021). Taking these first two assumptions together, cognitive constraints and biases make it very hard to anticipate how much time will be required to reach the tipping point and realize the value of a novel recombination.

This timing prediction problem becomes dangerous, if not catastrophic, for firms under a third assumption of inefficiencies in the market for financial capital. If entrepreneurs and investors underestimate the amount of execution time needed, entrepreneurs will raise insufficient funding and investors will under-allocate funding needed for the firm to survive through the amount of execution time actually needed. This would be fine if capital markets were efficient and entrepreneurs could instantaneously get funding from investors. However, mobilizing additional resources from investors for a new funding round is inefficient and can take many months to even years. Nanda & Rhodes-Kropf (2016) outline this “financing risk” where even startups that are fundamentally sound can fail to mobilize the necessary resources in time. Therefore, assuming impaired foresight under complexity and inefficient resource allocation, the complexity inherent in novel recombination can lead to increased incidence of firms running out of time and being unable to subsequently secure additional funding in the traditional manner from new investors. Instead, these firms must either rely on bailout bridge financing from existing investors or cease operations altogether.

Hypothesis 5 *Pursuing more novel recombination increases a startup’s likelihood of requiring unanticipated additional resources and also of shutting down.*

In the context of venture-backed entrepreneurship, unanticipated needs for additional resources cause two scenarios which can be empirically observed. First, if a startup fails to meet the unexpected resource requirements, it may be forced to shut down. In a reductive sense, entrepreneurial failure is simply running out of money. As long as the entrepreneur can continue raising funds, the company is not considered to have failed. However, when the time requirements exceed what investors are willing to bear and what entrepreneurs can practically justify persevering through, startups choose to or have to shut down.

The second scenario caused by unanticipated needs for additional resources is emergency bridge financing. In bridge financing, entrepreneurs approach existing investors to secure additional capi-

tal beyond what was initially provided. This bridge financing aims to “bridge the gap” between the originally predicted time requirements (which were underestimated) and the actual time needed to achieve a milestone and raise another formal funding round with new investors. Bridge financings are crucial to understanding complexity in new ventures but have not been extensively studied in the existing literature. By considering bridge rounds, we can capture the challenges posed by complexity that would have led to shutdowns if not for the bailout from existing investors. To the best of our knowledge, the bridge financing phenomenon is overlooked in the management literature and we are the first to draw attention to its importance.⁶

The example of SoftBank and Katerra helps illustrate the relevance of the bridge financing phenomenon for studying complexity and the need for unanticipated additional resources. In January 2018, SoftBank invested \$865M in the Series D round of Katerra, a technology company vertically integrating and modularizing construction founded by the former CEO of Flextronics. Despite seemingly relevant experience with the integrated modular strategy, there were unforeseen interdependencies in the construction context. Katerra experienced delays and cost overruns leading to an increased burn rate and an emergency need for additional capital. SoftBank provided a \$200M bridge financing to allow Katerra to continue to work towards its vision. When the complex context again led Katerra to fail to reach the milestone in the timeframe, Softbank provided an additional \$200M bridge financing. Finally, a year later, Katerra shut down and went bankrupt. SoftBank’s bridges proved to be bridges to nowhere.

Katerra’s story is not unique. To the contrary, requiring bridge financings (that often end up being bridges to nowhere) is a common occurrence. However, the antecedents to a need for a bridge financing—rather than the choice of the investor to allocate a bridge financing—have not been system-

⁶Even in the finance and accounting literature, relatively few studies address bridge financing, with the notable exception of [Ewens et al. \(2016\)](#), who study the theory of hold-up and show that bridge financings are more likely to fail and less likely to lead to IPO. We do not observe any studies that use the need for bridge financing as an outcome variable to measure firm performance.

atically studied in the management literature. While we use bridge financings to measure slower-than-anticipated growth and larger-than-anticipated resource requirements in this study, they represent an important context that warrants further investigation, as we will detail in the *Discussion* section.

2.4 EMPIRICAL METHODS

2.4.1 CONTEXT

Venture-backed growth startups provide an ideal setting to explore our theory and test our hypotheses. The entrepreneurs in this setting are often at the frontier of technological progress and innovation. As such, recombination is widespread, readily observable, and likely to contain instances of novel recombination. In this setting, there are readily available proxies for our key theoretical constructs and hypothesized relationships: information on the opportunities and recombinations, the timing of execution, and measures of performance.

2.4.2 DATA AND SAMPLE

Our primary dataset consists of Pitchbook venture data combined with LinkedIn employment data and BuiltWith product architecture data. Pitchbook is an aggregator of data on private capital markets, including new ventures and venture capital investment in those ventures. Pitchbook uses a combination of technical software tools and a large human data operations team to assemble what is arguably the largest and most accurate database of startups and venture capital investment in those startups. Relevant to our interests, Pitchbook contains data on all the dimensions necessary to test our hypotheses including (i) the nature of each opportunity with detailed descriptions of the opportunity pursued by each startup, (ii) the amount of time required to reach key milestones, and (iii) the performance of the startup over time.

We construct a sample of entrepreneurial firms from Pitchbook to test our hypotheses. We start

with deal-level data from the Pitchbook database as of January 2022. We only keep completed deals, excluding those deals that are currently in process, that lack a date, or that fell through such that our analysis only examines deals that actually took place. Further, we also restrict our sample to include only deals of firms that raised at least one round that was a “Series” (e.g., “Series A”) in order to ensure that all firms in our sample are high-growth-intended rather than purely subsistence-focused (Hurst & Pugsley, 2011).

Given that one of our key independent variables of interest is a measure of novelty, we restrict our sample to firms founded from the year 1996 through 2020, although we use the full set of entrepreneurial firms to create that measure of novelty. The first time a combination occurs in the dataset, every combination of that firm is considered novel by construction even if it already occurred previously (because no prior combinations are captured in the dataset). Consequently, by restricting our data to firms founded from the year 1996 onward, we can have confidence that a combination that appears as novel truly is so. Furthermore, we exclude firms founded after 1996 given that this would not practically leave a reasonable amount of time to see outcomes or even subsequent funding rounds in many cases. In addition, we do not include funding rounds beyond 10 rounds, given that most startups raise far fewer funding rounds and therefore going beyond 10 (e.g. to a “Series K”) signals that something unusual is happening that should be excluded. These criteria result in a sample of 31,450 entrepreneurial firms across 94,549 funding rounds.

In order to proxy an organization’s need to manage complexity, we also collect data on the organizational structure and activities of firms over time. We match our Pitchbook sample to the LinkedIn database using the firm webpage URLs.⁷ In particular, LinkedIn contains data on the employees working at a given company over time, including the descriptions of the activities performed by employees. As asserted in Section 2.3.1, we argue that the activities performed by employees can serve as a reasonable proxy for the complexity and interdependencies the organization is seeking to manage. As

⁷Appendix Section B.1.2 outlines this match process and sample construction.

a result, as the level of novel recombination increases, we expect to see an increase in the the activities relevant to managing complexity and interdependencies.

In addition, we conduct two separate analyses using the BuiltWith data, contained in Appendix Section B.1.1 and Appendix Section B.1.5. BuiltWith data contains detailed information on the construction of webpages over time. For software firms where the product is the webpage, we can use this data to construct measures about the product over time.

We combine this data to suit the empirical needs of each hypothesis, i.e., firm-year level (*Hypothesis 6*) and firm-round level (for *Hypotheses 7* and *8*).

2.4.3 VARIABLES

DEPENDENT VARIABLES

To assess *Hypothesis 6*, we construct *Manage Complexity* and *Manage Third Parties* at the firm-year level.

MANAGE COMPLEXITY *Manage Complexity* is equal to the total number of words in the job activity descriptions of all employees working at the firm in a given year that refer to managing general interdependencies.⁸ *Manage Complexity (Count)* is equal to the number of employees in a given year who have any of the relevant words in their job description.

MANAGE THIRD PARTIES One type of complexity that firms face is managing the interdependencies with various third parties (e.g., suppliers, complementors, alliances). *Manage Third Parties* is equal to the total number of words in an employee's job activity description that refer to managing third parties.⁹ *Manage Third Parties (Count)* is equal to the number of employees in a given year who

⁸The complete list of words can be found in Appendix Section B.1.3.

⁹The complete list of words can be found in Appendix Section B.1.3.

have any of the relevant words in their job description.

For *Hypothesis 7*, we track a firm's valuation over time. Given that valuation and performance outcomes are not observed regularly each year, but rather are only observed at each funding round, we use the dataset at the firm-round level.

PREMONEY VALUATION *Valuation* is equal to the focal firm's premoney valuation (in millions of USD). This represents the valuation of the firm in the focal round excluding the new cash invested as part of the focal round. To derive this premoney valuation, we take the postmoney valuation provided by Pitchbook and subtract the amount invested in the focal round. We control for amount invested in the focal round separately.

POSTMONEY VALUATION *Valuation* is equal to the focal firm's postmoney valuation (in millions of USD). This represents the valuation of the firm in the focal round including the cash invested as part of the focal round. We still control for amount invested in the focal round separately.

For *Hypothesis 8*, we assess a firm's need for unforeseen resources through *Receive Bridge Financing*, *Need Bridge Financing*, and *Shutdown* at the firm-round level.

RECEIVE BRIDGE FINANCING We first identify bridge financings in the Pitchbook dataset. We define bridge financing as a financing labeled by Pitchbook as a bridge financing or where the type is the same as in the prior round (e.g., both "Series A"), over half the investors are follow-on investors, and it is a flat or a down round. This latter set captures the scenario where existing investors may choose to bridge the firm through equity rather than debt. *Receive Bridge Financing* is an indicator variable for whether the firm received a bridge financing immediately following the focal round.

NEED BRIDGE FINANCING However, many firms that run out of resources actually need a bridge financing but are unable to secure the bridge and are therefore forced to shut down. *Need Bridge*

Financing is an indicator variable for whether the firm raised a bridge financing or shut down immediately following the focal round.

SHUTDOWN *Shutdown* is an indicator variable for whether the focal firm shut down immediately following the focal round. To construct this measure, we start with the set of firms PitchBook labels as going bankrupt, liquidating, or going out of business. However, using this approach implies only a very small percentage of firms fail and shut down. But we know that most entrepreneurial endeavors fail. As is common in entrepreneurship datasets, failure is under-reported in the Pitchbook database. To correct for this under-reporting, we infer that a company has shut down if its website, as listed in the Pitchbook dataset, is either non-existent or fails to load.

However, even among the websites that load, it may simply load an error message or redirect to an irrelevant page (e.g., a domain of a failed company now for sale might redirect to a domain-purchasing site like GoDaddy). We examine the set of websites that load but that Pitchbook labels as failed, and we use this set of known false-positive webpages to identify the most common set of words associated with false positives (“found 404,” “forbidden 403,” “squarespace webpage expired,” etc.). We then include all webpages that include one of these sets of words associated with failure in the shutdown set. All in all, this updated measure suggests 36.2% of firms fail—a figure much closer to what might be expected.

While not formally hypothesized, we also assess the impact of novel recombination on a firm’s likelihood of achieving *IPO* at the firm-round level.

IPO *IPO* is an indicator variable for whether the firm went public via an initial public offering immediately following the focal round.

INDEPENDENT VARIABLES

NOVEL RECOMBINATION Following the definition of Fleming (2001), we construct a measure of “a new combination of components or a new relationship between previously combined components” (p. 118).¹⁰ To do this, we first assemble all the words used to describe the components of each entrepreneurial firm, including all the keywords and location, and convert them to word stems to consolidate tenses and variants of the same word. Taking a unique list of these words, we then exclude all common stop words, prepositions, numbers, and question words that do not represent relevant components of a firm. We also eliminate words that occur with great frequency and are therefore limited in their unique information value (e.g., company, business). In order to consolidate keywords that are synonyms, we collapse any keyword that appears together with another keyword in more than 10% of its total occurrences. We also eliminate firms where there is insufficient information in the keywords to describe the actual business activities by dropping firms with fewer than five keywords. We then compute all unique combinations of the remaining components for each firm. For example, a firm with the three components *electric*, *car*, and *Europe* has three unique combinations: (i) *electric-car*, (ii) *electric-Europe*, and (iii) *car-Europe*.

To determine the novelty of each combination, we sort each company by the year it was founded and compute the cumulative count of each combination by year (e.g., how many firms pursued a given combination in all the years prior to the focal year). Specifically, a recombination is considered novel if no other companies include that combination in any year prior to the year the focal company

¹⁰Much of the literature on recombination uses patents (Fleming & Sorenson, 2001, Fleming, 2001, Aggarwal et al., 2020) or research papers (Uzzi et al., 2013) as context. Although these settings allow for clear measurement, patents and research papers are often just abstract ideas and theoretical concepts that have not actually been executed, created, or implemented in practice. Therefore, considering only these abstract contexts may only allow for a partial picture of the implications of novel recombination. We assert that novel recombination in the abstract leads to complexity in the concrete process of execution, which can be costly to manage. Overlooking the complexity inherent in executing novel recombination drastically undervalues the actual hard work and resources required by the entrepreneur to realize abstract novel recombination in practice.

was founded. We restrict our sample to 1996–2020. As noted in Section 4.2, given that the first instance of any keyword would lead to a novel combination by construction, we define a pre-period from 1800–1995 and use the corpus of combinations in this pre-period as a baseline, so any observed novel recombination post-1995 would not simply be novel by construction.¹¹ To define the set of eligible keywords, we include all the keywords of companies founded in the pre-period up through 1995, and we eliminate all keywords in our sample of firms founded after 1995 that do not appear in this pre-period.¹² The *Novel Recombination* variable is the total number of combinations for each firm that do not appear in any year prior to the focal year. *Novel Recombination Binary* is an indicator variable equal to one if *Novel Recombination* is greater than or equal to one and equal to zero otherwise.

Given the importance of this measure, we conduct several additional analyses to validate that *Novel Recombination* captures the intended theoretical construct. We compare our measure of *Novel Recombination* to other proxies of novelty. First, we find a strong and positive correlation between *Novel Recombination* and a firm being among the first to enter a market and pursue a given opportunity. This relationship holds across four different methods of defining opportunity clusters on which to base entry order. Second, for the subset of firms where the product is a website, we find a positive correlation between *Novel Recombination* and novel combinations in the tech stack used to build the product. Taken together, these robustness checks show the relationships that would be expected when *Novel Recombination* captures the joining of two heretofore disparate components. Appendix B.1.1 describes these additional analyses in detail.

¹¹Starting after 1995 allows us to include most dotcom companies in our sample, the majority of which were founded in the late 1990s.

¹²Defining keywords based on the pre-period also helps account for any one-off variations in how terms are spelled given that it requires all keywords to occur at least two times.

FIRM AGE We define *Firm Age* as the the time elapsed since founding: the difference (in years) between the date of the focal round and the founding date of the focal firm.¹³

2.4.4 CONTROL VARIABLES

To account for the potential difference in impact of *Novel Recombination* in systems with varying firm scope, we control for *Firm Scope*: the total number of components for each firm, e.g., in our prior example of a firm with the three components (i) electric, (ii) car, and (iii) Europe, *Firm Scope* would be equal to three. To account for the differential impact of how much money a firm raises in a given round and cumulatively, we include the controls *Round Size* and *Cumulative Funding*, respectively. *CEO Change* is a binary variable equal to one if the CEO at the time of the focal round is different from the CEO at the time of the prior round; this controls for changes in management style and direction over time.

For regressions using LinkedIn employment data, we also include controls for *Total Employees* (to account for changing firm size), *Total Descriptions* (to account for changing numbers of employees with job descriptions), and *Description Words* (to account for variation in the number of words and associated detail provided in the descriptions).

In addition, as will be discussed subsequently in Section 3.3, we also include many fixed effects to control for unobserved heterogeneity correlated with our variables of interest. These include controlling for the funding-round number, the year of the deal, the firm's founding year, the firm's headquarters country, the firm's size, the firm's industry, and, in select models, the firm itself.

¹³Given that the funding round date has a day, month, and year whereas the founding date only includes the year, we assume January 1 for the month and year of founding. In the absence of a founding date, we use the year of the firm's first funding round.

2.4.5 DATA STRUCTURE AND MODEL

FIRM-YEAR LEVEL

Table 2.1 contains summary statistics for the variables at the firm-year level. Here, LinkedIn contains granular, continuous data on employees, including information on which employees worked at each company at any given point in time and the tasks they performed. As such, we are able to construct a panel of the activities managed by employees in each firm each year.

Table 2.1: Summary Statistics - Firm-Year Level. 114,598 firm-year level observations.

Statistic	Mean	St. Dev.	Min	Max
Manage Complexity	16.76	87.53	0	5,646
Manage Complexity (Count)	9.97	47.90	0	3,264
Manage Third Parties	26.64	194.82	0	20,778
Manage Third Parties (Count)	12.68	73.70	0	6,287
Novel Recombination	3.24	5.94	0	188
Firm Age	6.78	4.72	0	24
Firm Scope	7.12	2.24	5	34
Cumulative Funding	42.38	383.99	0	40,197.56
Round Size	11.00	185.94	0	39,000.13
CEO Change	0.07	0.26	0	1
Total Employees	56.64	296.41	1	28,311
Total Descriptions	24.18	120.45	1	10,003
Description Words	971.47	4,815.57	1	378,152

We estimate ordinary least squares (OLS) models across all analyses using this firm-year level panel. These models include fixed effects to control for several dimensions of otherwise unobserved heterogeneity that could correlate with the observed independent variables. *Year* fixed effects control for any variations or influences specific to a given year (e.g., certain terms became more widely used over time). *Round* fixed effects account for the concern that certain outcomes may be more likely to occur earlier on for lower-round numbers (e.g., a firm is less likely to go public on its first funding round but more likely to go public after raising many funding rounds). *Country* fixed effects control for potential heterogeneity across different geographies. *Industry* fixed effects take as a baseline firms that

involve technology but control for situations where the outcome of the firm largely does not involve technology or where the value proposition is mainly related to non-technical factors (e.g., regulation and macroeconomic conditions). Industry categories include firms in financial services, healthcare services, energy, and materials and resources. Given that this data is at the funding-year level, in order to account for potential correlation in the error term across years we generally cluster robust standard errors at the year level.

FIRM-ROUND LEVEL

Table 2.2 contains summary statistics for the variables at the firm-round level. Because we are studying private firms, performance data between funding rounds (e.g., every year on a continuous basis) is generally unavailable. As such, constructing a firm-year panel of performance would be impossible. We use this firm-round structure when looking at firm performance because it still provides a valuable snapshot of firm performance at each funding round. Also note the smaller sample size for the *Valuation* variable. This is due to the reality that interim valuation data is only available for a subset of funding rounds.

We estimate ordinary least squares (OLS) models across all analyses using this firm-round level panel. Similar to the prior model, we also include *Round*, *Country*, and *Industry* fixed effects. However, because our data is not at the year level, we do not include year fixed effects. Instead, we include two additional fixed effects that attempt to control for time-variant heterogeneity. *Round Year* fixed effects control for any factors that change year to year but are common to all funding deals between entrepreneurial firms and investors in a given year. *Founding Year* fixed effects control for any factors that change year to year but are common to all entrepreneurial firms launched in a given year. In our models testing how value increases over time, we include *Company* fixed effects, which allows us to compare within-firm variation in value over time by controlling for all unobserved heterogeneity at the firm level, including variation in industry, novelty, etc. Given that this data is at the funding-round

Table 2.2: Summary Statistics - Firm-Round Level. 94,549 firm-round level observations. Given that bridges are generally not reported when they occur immediately preceding shutdown, there are 73,212 observations where observing Receive Bridge Financing is possible. Given that valuation data is not always shared in private transactions, there are 52,283 recorded observations for Premoney Valuation and 47,657 observations for Postmoney Valuation.

Statistic	Mean	St. Dev.	Min	Max
Receive Bridge Financing	0.11	0.31	0	1
Need Bridge Financing	0.20	0.40	0	1
Shutdown	0.12	0.33	0	1
IPO	0.01	0.11	0	1
Premoney Valuation	168.26	1,485.65	0	178,000
Postmoney Valuation	213.07	1,695.04	0	180,000
Novel Recombination	2.90	5.71	0	188
Novel Recombination Binary	0.60	0.49	0	1
Firm Scope	7.20	2.24	5	34
Round Size	24.31	185.46	0	39,000.13
Cumulative Funding	42.16	339.87	0	40,197.56
Firm Age	4.98	3.67	0	25.98
CEO Change	0.14	0.34	0	1
Round	2.99	1.92	1	10
Round Year	2,015.24	4.89	1,996	2,021
Founding Year	2,010.76	5.55	1,996	2,020

level, in order to account for potential correlation in the error term across firms of the same round number we generally cluster robust standard errors at the funding-round level.¹⁴

2.5 RESULTS

Hypothesis 6 predicts that novel recombination leads to an increasing need to manage complexity. We find support for this hypothesis by looking at the relationship between novel combinations and the organizational activities of the entrepreneurial firm. We examine the relationship between novel recombination and the number of words related to managing complexity and interdependencies across all the employee's job activity descriptions in a firm each year. As outlined in Table 3.5, we find that *Ln Novel Recombination* is positively correlated with *Manage Complexity* ($p \approx 0.000$) and with *Manage Complexity (Count)* ($p \approx 0.000$). Doubling the total *Novel Recombination* increases the number of words related to managing complexity in an employees' job activity description by approximately 3.13 words and increases the number of employees with any reference to managing complexity by approximately 1.74 (representing increases of 18.7% and 17.5%, respectively, over the sample mean).¹⁵ We also find that *Ln Novel Recombination* is positively correlated with *Manage Third Parties* ($p \approx 0.000$) and with *Manage Third Parties (Count)* ($p \approx 0.000$). Doubling the total *Novel Recombination* increases the number of words related to managing third parties in an employee's job activity description by approximately 6.42 words and increases the number of employees with any reference to managing third parties by approximately 2.46 (representing increases of 24.1% and 19.4%, respectively, over the sample mean).

Appendix Section B.1.5 outlines an additional robustness check of this relationship between

¹⁴Appendix Section B.1.4 presents correlation tables for the data at both the firm-year level and the firm-round level.

¹⁵To interpret effect sizes based on logged independent variables here and elsewhere, we calculate the effect based on a doubling, or 100% increase, in the independent variable. Thus, in the case of the *Manage Complexity* independent variable, we take $4.515 * \ln(2) = 3.130$.

Table 2.3: Novel Recombination and Complexity. Ordinary least squares (OLS) estimation. Robust standard errors are clustered at the year level and shown in parentheses. *p*-values are shown in brackets.

	Manage Complexity	Manage Complexity (Count)	Manage Third Parties	Manage Third Parties (Count)
Ln Novel Recombination	4.515 (0.510) [0.000]	2.513 (0.288) [0.000]	9.261 (1.031) [0.000]	3.543 (0.405) [0.000]
Ln Novel Recombination x Firm Age	-1.010 (0.166) [0.000]	-0.544 (0.092) [0.000]	-2.102 (0.344) [0.000]	-0.785 (0.132) [0.000]
Firm Scope	4.501 (0.960) [0.000]	2.258 (0.498) [0.000]	8.889 (1.846) [0.000]	3.628 (0.762) [0.000]
Firm Age	1.544 (0.256) [0.000]	0.800 (0.143) [0.000]	3.293 (0.540) [0.000]	1.221 (0.210) [0.000]
Ln Cumulative Funding	2.851 (0.758) [0.000]	1.320 (0.382) [0.002]	6.099 (1.565) [0.000]	2.431 (0.611) [0.000]
Ln Round Size	-2.138 (0.456) [0.000]	-1.113 (0.240) [0.000]	-4.806 (0.927) [0.000]	-1.923 (0.363) [0.000]
CEO Change	-0.359 (0.664) [0.594]	-0.294 (0.403) [0.474]	0.514 (1.179) [0.667]	0.111 (0.498) [0.826]
Ln Total Employees	-9.032 (1.494) [0.000]	-4.975 (0.817) [0.000]	-14.898 (2.294) [0.000]	-6.254 (0.931) [0.000]
Ln Total Descriptions	59.130 (5.650) [0.000]	35.285 (3.315) [0.000]	106.883 (9.856) [0.000]	48.223 (4.286) [0.000]
Ln Description Words	-15.182 (1.830) [0.000]	-9.534 (1.137) [0.000]	-30.339 (3.592) [0.000]	-13.939 (1.631) [0.000]
Observations	114 598	114 598	114 598	114 598
R^2	0.196	0.225	0.123	0.173
Year	X	X	X	X
Round	X	X	X	X
Country	X	X	X	X
Industry	X	X	X	X

novel recombination and complexity. For the subset of software companies, we show a negative relationship between novel recombination and the opposite of interdependence: modularity. Modularity is proxied by API use, where the firm can easily plug in a modular, already-existing solution to perform part of the needed product function. We find that firms with higher levels of novel recombination use fewer modular components, which would be less likely to be compatible with more complex products that have more novel interdependencies.

Hypothesis 7 predicts that firms with more novel recombination will take a longer time to achieve the same valuation. We find support for this hypothesis by looking at the relationship between valuation and firm age for firms with any novel recombination versus those without novel recombination. As outlined in Table 2.4, we find that *Premoney Valuation* is positively correlated with *Firm Age* ($p = 0.005$) and negatively correlated with *Firm Age*² ($p = 0.174$). However, the opposite signed relationship is observed for the same terms of *Firm Age* ($p = 0.072$) and *Firm Age*² ($p = 0.067$) when interacted with *Novel Recombination Binary*. When *Novel Recombination Binary* equals one (signaling the firm has any novel recombination), the value of the base *Firm Age* term is slightly lower (subtracting 17.46 from the coefficient of 152.12 in the base case when *Novel Recombination Binary* is equal to zero). In addition, when *Novel Recombination Binary* is equal to one, the value of the squared term is slightly less negative (adding 1.15 to the coefficient of -0.67 in the base case when *Novel Recombination Binary* is equal to zero). As elaborated on in Section 2.5.1 and illustrated graphically in Figure 2.2, the net effect of these changes is that, when *Novel Recombination Binary* is equal to one, for smaller values of *Firm Age*, *Valuation* is lower compared to firms where *Novel Recombination Binary* is equal to zero. However, while not formally hypothesized, we do see that for larger values of *Firm Age*, the lines cross and *Valuation* is higher for firms where *Novel Recombination Binary* is equal to one.

As shown in Table 2.4, we run the same regression but instead use the dependent variable *Postmoney Valuation*. The directionality and significance of the results hold; we just see larger effect sizes, which makes sense given that postmoney valuations are larger due to counting the cash invested as

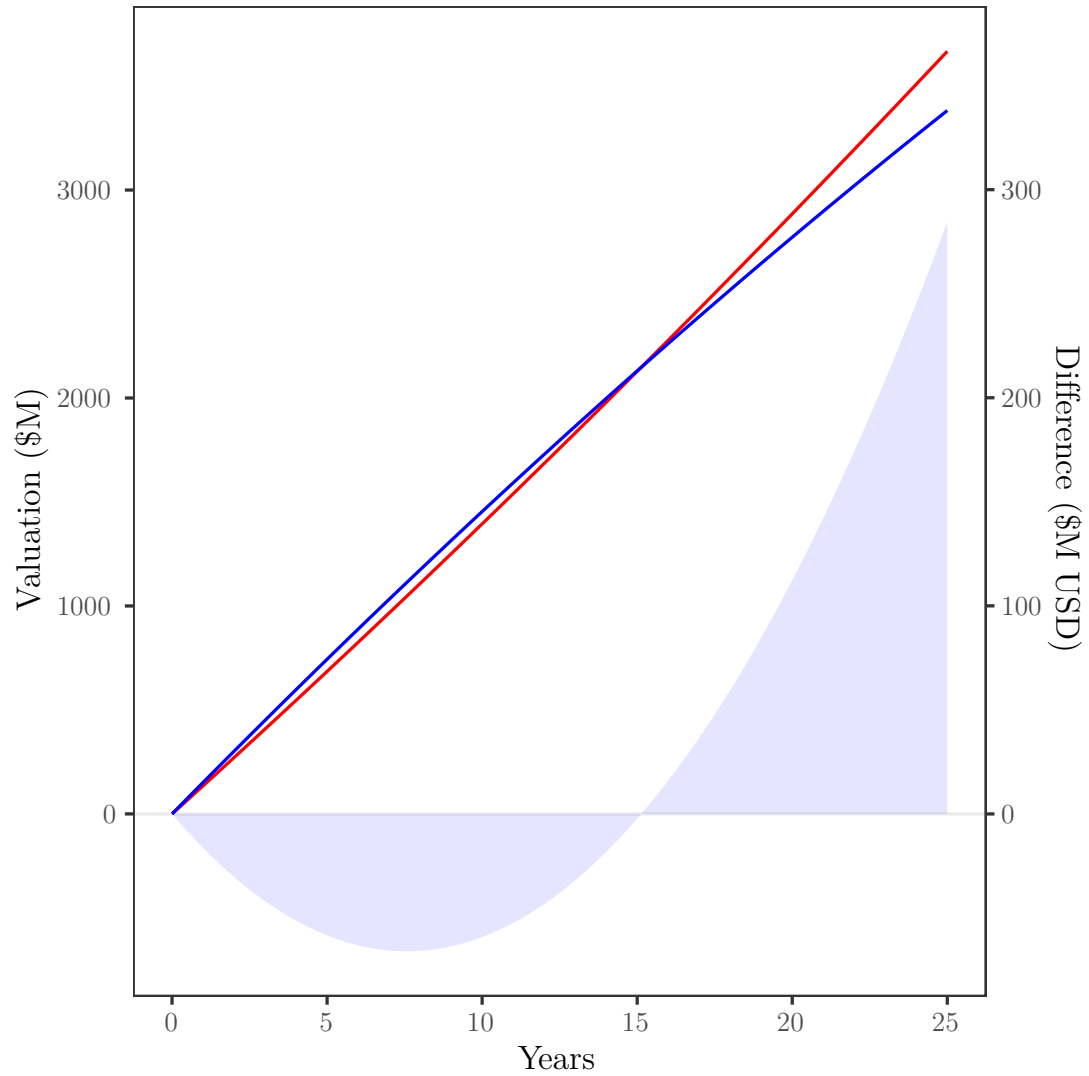


Figure 2.2: Novel Recombination and Valuation Over Time. The first vertical axis *Valuation* represents the pre-money valuation of the entrepreneurial firm. The second vertical axis *Difference* represents the difference between the two lines. The horizontal axis *Years* represents the total number of years since founding.

Table 2.4: Novel Recombination and Valuation Over Time. Ordinary least squares (OLS) estimation. Robust standard errors (in parentheses) are clustered at the round year level. p -values are shown in brackets.

	Premoney Valuation	Postmoney Valuation
Firm Age	152.121 (49.553) [0.005]	185.616 (62.460) [0.006]
Firm Age ²	-0.673 (0.481) [0.174]	0.044 (0.652) [0.947]
Novel Recombination Binary * Firm Age	-17.460 (9.299) [0.072]	-19.647 (8.824) [0.035]
Novel Recombination Binary * Firm Age ²	1.154 (0.603) [0.067]	1.349 (0.708) [0.068]
Observations	52 283	47 657
R^2	0.547	0.545
FE: Round Year	X	X
FE: Firm	X	X

part of the focal round.

To further explore the relationship between novel recombination and time to realize value, we explore how novel recombination influences the time it takes a firm to clear the \$100M and \$1B valuation marks. We subset the data to the 4,524 firms that ever reach a valuation of \$100M (or greater) and find that, among those firms, doubling *Novel Recombination* leads to a 1.7 month increase in *Time to \$100M Valuation* ($p \approx 0.000$). This represents a 2.3% increase over the sample mean of 6.2 years. We then subset the data to the 844 firms that ever reach a valuation of \$1B (or greater) and find that, among those firms, doubling *Novel Recombination* leads to a 4.3 month increase in *Time to \$1B Valuation* ($p = 0.005$). This represents a 4.7% increase over the sample mean of 7.6 years.

Hypothesis 8 predicts that, as the number of novel combinations increases, the likelihood of needing unanticipated additional resources and shutting down also increases. We find support for this hypothesis by looking at the relationship between novel recombination and needing bridge financing, receiving bridge financing, and firm shutdown. In Table 2.6, we find that *Novel Recombination*

Table 2.5: Novel Recombination and Time to High Valuation. Ordinary least squares (OLS) estimation. Robust standard errors (in parentheses) are clustered at the year-founded level. p -values are shown in brackets.

	Time to \$100M Valuation	Time to \$1B Valuation
Ln Novel Recombination	0.205 (0.050) [0.000]	0.512 (0.166) [0.005]
Firm Scope	0.046 (0.025) [0.077]	-0.067 (0.036) [0.072]
Ln Round Size	-0.201 (0.098) [0.052]	0.195 (0.157) [0.225]
Ln Cumulative Funding	0.604 (0.111) [0.000]	-0.213 (0.168) [0.217]
CEO Change	-0.023 (0.107) [0.830]	-0.375 (0.304) [0.229]
Observations	4524	844
R^2	0.454	0.648
FE: Founding Year	X	X
FE: Country	X	X
FE: Industry	X	X

positively correlates with *Need Bridge Financing* ($p \approx 0.000$); doubling the number of novel recombinations increases the likelihood of needing a bridge subsequent to the focal round by 0.5 percentage points (a 2.5% increase over the sample mean). *Novel Recombination* positively correlates with *Receive Bridge Financing* ($p = 0.027$); doubling the number of novel recombinations increases the likelihood of receiving a bridge subsequent to the focal round by 0.2 percentage points (a 1.9% increase over the sample mean). *Novel Recombination* also positively correlates with *Shutdown* ($p \approx 0.000$); doubling the number of novel recombinations increases the likelihood of shutting down subsequent to the focal round by 0.4 percentage points (a 3.5% increase over the sample mean).

Table 2.6: Novel Recombination and Performance. Ordinary least squares (OLS) estimation. Robust standard errors (in parentheses) are clustered at the funding round level. p -values are shown in brackets.

	Need Bridge Financing	Receive Bridge Financing	Shutdown	IPO
Ln Novel Recombination	0.007 (0.0007) [0.000]	0.003 (0.001) [0.027]	0.006 (0.001) [0.000]	0.001 (0.0007) [0.120]
Firm Scope	-0.010 (0.001) [0.000]	-0.0004 (0.0007) [0.568]	-0.011 (0.001) [0.000]	-0.0008 (0.0002) [0.012]
Ln Round Size	0.004 (0.003) [0.228]	-0.015 (0.003) [0.000]	0.022 (0.003) [0.000]	-0.0006 (0.0009) [0.541]
Ln Cumulative Funding	-0.007 (0.006) [0.260]	0.013 (0.003) [0.001]	-0.026 (0.005) [0.000]	0.012 (0.001) [0.000]
Firm Age	0.010 (0.005) [0.107]	0.005 (0.005) [0.307]	0.016 (0.003) [0.001]	-0.003 (0.001) [0.018]
CEO Change	0.025 (0.003) [0.000]	0.013 (0.003) [0.003]	0.017 (0.003) [0.000]	-0.001 (0.001) [0.464]
Observations	94 549	73 212	94 549	94 549
R^2	0.031	0.021	0.048	0.031
FE: Round	X	X	X	X
FE: Round Year	X	X	X	X
FE: Founding Year	X	X	X	X
FE: Country	X	X	X	X
FE: Industry	X	X	X	X

As a matter of empirical exploration, we examine the relationship between novel recombinations and IPO. As outlined in Table 2.6, we find that *Ln Novel Recombination* is positively correlated with *IPO* ($p = 0.120$). Doubling *Novel Recombination* increases the likelihood of IPO on the subsequent round by 0.1 percentage points. Given that IPOs are relatively rare, this represents a 7% increase over the sample mean.

2.5.1 VISUAL EVIDENCE FOR THE TIME DIFFERENCE IN VALUATION

To illustrate the results of the model outlined in Table 2.4 more clearly, Figure 2.2 plots the relationship between *Premoney Valuation* and *Firm Age* in the case when *Novel Recombination Binary* is equal to zero and in the case when *Novel Recombination Binary* is equal to one. This allows us to clearly compare the evolution of value over time for firms pursuing novel recombination versus the evolution of value over time for firms not pursuing novel recombination. Initially, over the first 12 years post-founding, the valuation of firms pursuing novel recombination lags behind that of firms that are not. There is as much as a \$50M USD lag in value for firms pursuing any novel recombination. Figure 2.2 is reminiscent of the timing of value realization theorized in Section 2.3.2 and visualized in Figure 3.3.

While we do not hypothesize this, one might expect novel recombination to lead to higher potential value, explaining why in equilibrium entrepreneurs would still pursue these opportunities even if all our other observations around timing and failure are correct. We see that around a *Firm Age* of 15, the lines cross and at 25 years post-founding the valuation of firms pursuing novel recombination is almost \$300M USD greater compared to firms not pursuing novel recombination. This is consistent with the IPO evidence mentioned earlier, where firms pursuing more novel recombination appear to be more likely to IPO.

To illustrate the difference in time more clearly, Figure 2.3 plots the additional time required for a firm with any novel recombination to reach a given valuation. To calculate this difference, we use the model estimates from the first column of Table 2.4 to estimate the time required for a firm without

any novel recombination to reach a given valuation. We then calculate the amount of time it would take a firm with any novel recombination to reach the same valuation. We plot the difference between these two values. It takes firms pursuing any novel recombination about one month longer to achieve a \$100M USD valuation and almost six months longer to achieve a \$1B USD valuation.

2.6 DISCUSSION

We now highlight three contributions of this work. First, our complexity-based view of entrepreneurship recognizes the real challenge of resolving interdependencies when executing in practice, with implications for research on learning and experimentation. Second, we document a critical driver of timing at the firm level and the implications for the timing of value realization. Third, we highlight the importance of understanding entrepreneurial failure and put forth bridge financings as an important but overlooked measure of failure.

2.6.1 A COMPLEXITY-BASED VIEW OF ENTREPRENEURSHIP

A large body of work asserts that the ability to bear and eventually resolve uncertainty is a hallmark of entrepreneurship. Foundational work by Knight (1921) outlines uncertainty as the root of entrepreneurial profit, with the key task of the entrepreneur in unlocking that profit being to resolve and reduce the overall uncertainty. As such, a tolerance for uncertainty and an ability to resolve it are highlighted as prerequisites for individuals who choose to pursue entrepreneurship (Knight, 1921, Koudstaal et al., 2016) and are also associated with superior performance among those who do select into entrepreneurship (Levine & Rubinstein, 2017). However, we agree with Townsend et al. (2018a) that uncertainty is a broad, catch-all term that is imprecisely used to describe what are in fact many distinct knowledge problems.

In particular, complexity is often conflated with uncertainty despite being a distinct knowledge




Figure 2.3: Novel Recombination and Time to Catch-up. The horizontal axis represents the valuation of a firm without any novel recombination. The vertical axis represents the total time (in months) it would take the firm pursuing any novel recombination to catch-up and achieve a comparable valuation.

problem that, as we will discuss next, warrants separate study. As noted previously, problems of uncertainty are driven by an unknown variable. But this variable is a known unknown (e.g., the entrepreneur has a specific and known question to which she does not know the answer). For example, an entrepreneur might need to predict whether there is demand for nightly rentals in a particular market. On the other hand, complexity stems from the multiplicity of variables and interactions between them, involving coordination between components and/or firm activities (Simon, 1962, Porter & Siggelkow, 2008). In this context, there are many unknown-unknowns (e.g., things that the entrepreneur does not know she does not know ex ante and that are only knowable upon executing). For example, even in the absence of pure uncertainty, knowing for certain that there is demand for nightly rentals and that there are people with empty beds who want to collect money from nightly renters, an entrepreneur might not know whether both can be served simultaneously because there are unknown interdependencies between these groups that will only become known and, as a result, can only be addressed during execution.

Given that resolving problems of complexity requires a different approach from resolving problems of uncertainty, we argue a complexity-based view as opposed to an uncertainty-based view that has important implications for entrepreneurial strategy. Importantly, problems of uncertainty can be resolved by obtaining additional information to fill the known gap: “The actor knows the question being asked...thus, the discovery of critical data through entrepreneurial action will resolve the knowledge problem” (Townsend et al., 2018a, p. 675). Learning, experiential or vicarious, resolves uncertainty. The extant literature demonstrates that additional information gained through experience decreases uncertainty and thereby increases performance and value over time (Rapping, 1965, Darr et al., 1995). More recent work demonstrates these same principles in the entrepreneurial context (Gans et al., 2019, Politis, 2005). In contrast, complexity does not lend itself to resolution through traditional learning models. In complex systems the relationship between inputs and outputs may not be known, which can severely limit the value of additional information given that it is unclear which

inputs are actually driving the desired outputs.

In some instances, additional information can even decrease performance in complex systems. Given that bounded rationality places an upper bound on the information an entrepreneur can reasonably process, complexity can lead entrepreneurs to come away with the wrong learnings, thereby decreasing performance or even leading to increased failure (Eggers & Song, 2015, Hallen & Pahnke, 2016). In fact, as complexity increases, learning is shown to decrease entrepreneurial foresight (Peterson & Wu, 2021). Further research should explore the performance of traditional learning models under complexity.

IMPLICATIONS FOR ENTREPRENEURIAL EXPERIMENTATION

The implications of a complexity-based view of entrepreneurship place important caveats on the increasingly popular notion of entrepreneurial experimentation (Camuffo et al., 2020), which we can think of as a strategy of controlled, intentional learning. A/B testing is a specific experimental approach used to reduce uncertainty, and the recent surge in articles on this topic attests to its prevalence among entrepreneurs (Koning et al., 2022, Bhat et al., 2020). Indeed, the popular Lean Startup approach and accompanying Silicon Valley mantra to “fail fast” encapsulate this push to reduce uncertainty (even if the outcome is failure) through rapid experimentation (Reis, 2011). Facebook’s infamous motto to “move fast and break things” also reflects this tendency.¹⁶ However, given the limitations of complexity on learning—limiting an entrepreneur’s ability to relate inputs and outputs—it of

¹⁶Under this view, reducing the uncertainty inherent in entrepreneurial opportunities through experimentation is the foundation of the tranced nature of venture capital investment (Kerr et al., 2014). While the expected value of an entire opportunity may be negative given the inherent uncertainty, the venture capital financing process breaks the opportunity down into a series of experiments which, if successful, decrease the uncertainty of the overall opportunity. This breakdown creates tranced investment rounds (e.g., Series A, Series B), which can have positive expected value in isolation. Given that this calculation of risk and uncertainty determines venture capital funding, the total resources available to a startup over a given period of time are determined by the degree of expected resolution in uncertainty. Additional information is gathered across each financing round through the efforts enabled by the invested resources which, if successful, then decrease uncertainty and increase the value of the startup.

course follows that experimentation also does not resolve complexity and therefore is of limited value in complex systems (Townsend et al., 2018a). In complex systems, entrepreneurs must confront unknown unknowns where the specific challenge and problem are not clear ex ante. Thus, hypothesizing and then experimenting in complex systems may be precluded entirely given that the entrepreneur is not even aware of the relevant interdependencies. Another way to conceptualize this may be that in complex systems the costs of experimentation can be preclusively high.

2.6.2 TIMING

By drawing attention to the importance and determinants of timing at a firm level, we draw attention to overlooked implications of knowledge recombination and industry-level S-curves.

Much of the existing literature on innovation and knowledge recombination makes the assumption that innovation is instantaneous and there is no time requirement to realize value after recombination. As noted, patents are one of the most common measures of novel recombination and the value of patents is generally assumed to be inherent and therefore immediate (Fleming & Sorenson, 2001, Fleming, 2001, Aggarwal et al., 2020). However, we argue that while using patents as a dependent variable does measure novel recombination in the abstract, it is not a measure of the actual realization of value in the market for products and services. An opportunity's value is not realized until after the time required for execution is invested in practice. As alluded to previously, the first patents on electric vehicles appeared in the late 1800s. However, it has taken centuries of time for the recombination to be executed in practice and for a critical number of the relevant interdependencies to be resolved. Indeed, even now, the complexities of this recombination continue to slow and delay the realization of the full potential of the electric vehicle opportunity.

The related literature on industry evolution does consider time (Utterback & Suárez, 1993, Suarez & Lanzolla, 2007); indeed, time is generally the x-axis on the classic S-curve plot. Even more relevant to this study, Adner & Kapoor (2016) specifically study the “pace” of jumps between curves; one

could consider the novel opportunities pursued by entrepreneurs in our study as representing jumps to a new curve. However, the S-curve is generally intended to reflect progress at the industry level. In contrast, we are concerned with the individual firm, and the individual firm does not benefit from the law of averages implicitly represented in an industry-level S-curve. In other words, the industry is concerned with when the shift to a new standard or technology will be realized, but the individual firm is concerned with whether it will be the one to execute on its pursued opportunity on a practical timeline. For example, consider the evolution of the pet retailing business as it shifted to online retailing. The early Pets.com is a well known “failure” of the dot-com era—or, from our perspective, it simply ran out of time. It very well could have succeeded and captured all the value eventually realized by Chewy two decades later. This same story of a pioneering entrant running out of time has played out many times, e.g., grocery delivery introduced by Webvan (versus Amazon Fresh), personal digital assistants introduced by Apple’s Newton (versus PalmPilot). The point is that Webvan and the Newton were fundamentally good ideas, but they just took more time to realize than proved practical at the time. If it takes too long to address the relevant interdependencies, the firm may not be able to persevere through execution to achieve the value of an otherwise good opportunity.

2.6.3 ENTREPRENEURIAL FAILURE

A significant portion of entrepreneurship literature primarily focuses on success (e.g., [Stuart et al., 1999](#), [Kroll et al., 2007](#), [Kaplan et al., 2009](#), [Dahl & Sorenson, 2012](#)), while a smaller portion delves into failure (e.g., [Artinger & Powell, 2016](#), [Eggers & Song, 2015](#)). However, examining failure in entrepreneurship is crucial as it can provide meaningful insights into strategy. While Michael Porter’s work centered on understanding why companies succeed ([Porter, 1985](#)), Clayton Christensen’s research originated from an interest in why firms fail ([Christensen, 1997](#)). By focusing on the determinants of failure, Christensen was able to offer unique and important insights that may not have been possible with a sole focus on success. While existing literature focuses on the emergence of high-

performing returns from novel recombination, in fact, looking through the lens of failure tells a very different story. Because novel recombination delays the realization of the opportunity, through the mechanism of complexity, the risk of failure is actually increased. Therefore, for an entrepreneur seeking to minimize failure rather than maximize success, novel recombination may not be as attractive. Considering failure specifically can provide this type of entrepreneur with insights that cannot be gained solely from studying success.

Bridge financing has received little to no attention in prior research; however, it contains important information about entrepreneurial failure. Bridge financing represents a specific type of failure experienced by entrepreneurs and venture capital investors. It occurs when entrepreneurs run out of resources before achieving the necessary milestones to raise a formal funding round and must seek a bridge financing from existing investors.

To understand why the existence of bridge financing contains important information about entrepreneurial failure, we have to consider the antecedents to the emergence of a bridge financing, both from the entrepreneur's and venture capitalist's perspectives. From the entrepreneur's perspective, the need for a bridge financing reflects a failure to secure the resources required to reach an intended milestone and raise a subsequent formal funding round. From the venture capital firm's perspective, when one of their investments requires a bridge financing this reflects a failure in the firm's assessment of the entrepreneur and the opportunity. Both from the perspective of the entrepreneur and from the perspective of the investor, only looking at companies that ultimately succeeded would miss part of the picture, notably any failures encountered along the way that required bridge financings to correct. Similarly, only looking at companies that ultimately failed also misses part of the picture, notably any companies that also failed but were able to secure a bridge financing and continue. In entrepreneurship, a company is never really a failure as long as it can keep raising money; bridge financings reflect those startups that persist despite failure.

The context of bridge financing presents numerous opportunities for further research, including

studying the success rate of bridge financings compared to traditional financing rounds, exploring the characteristics of firms that successfully secure bridge financings and survive, and investigating the types of entrepreneurs and venture capitalists more likely to require bridge financings.

2.7 CONCLUSION

This study addresses how novel recombination and entrepreneurial opportunity in the abstract require concrete execution and management of complexity and interdependencies to realize in practice. Importantly, complexity impacts how the value of an opportunity is realized over time. For opportunities with high levels of novel recombination and therefore high levels of complexity, the time required to execute is greater and, consequently, the realization of value is delayed. Importantly, when managing complexity, time requirements can be hard to anticipate such that unanticipated delays and even failure are more frequent.

3

The Mediocrity Trap in Entrepreneurial Learning

3.1 INTRODUCTION

Effective and efficient learning is a critical determinant of entrepreneurial performance (March, 1991, Minniti & Bygrave, 2001, Politis, 2005). New ventures face a high degree of uncertainty in trying to

identify valuable opportunities (Alvarez & Barney, 2007, Gruber et al., 2008) and often seek to resolve that uncertainty by learning from the market (Townsend et al., 2018b, Gans et al., 2019). This uncertainty primarily concerns two separable dimensions—product and market (Helfat & Lieberman, 2002, Moeen et al., 2020, Agarwal et al., 2004)—which together constitute product-market fit (Zellweger & Zenger, 2023). At the same time, new ventures are also generally resource-constrained (Baker & Nelson, 2005, Katila & Shane, 2005, Clough et al., 2019). Put simply, new ventures don't have a lot but also need to learn a lot. As a result, entrepreneurs need to try to learn and resolve uncertainty in a resource-efficient manner (Kerr et al., 2014, Felin et al., 2020). Entrepreneurs cannot afford to waste resources on building and trying to sell products no one actually wants to buy. As such, extant literature recommends that new ventures initially invest in minimal versions of their product to learn and validate market demand quickly so as not to persist and waste limited resources where demand does not exist (Camuffo et al., 2020). This approach has been widely embraced as a way to help entrepreneurs “fail fast” and avoid wasteful persistence.

But in practice, studies show new venture persistence and delayed exit, with firms continuing to pursue opportunities even after reasonable negative market signals, thereby wasting resources (Elfenbein et al., 2017, Chen et al., 2019). Even with the general popularity of prescriptions against persistence—not over-investing resources in developing a product and, instead, trying to learn by selling it quickly—new ventures continue to delay exit and persist (Chen et al., 2024). Some research points to behavioral biases, such as overconfidence (Camerer & Lovallo, 1999). Recent research also notes that persistence can be rational (Elfenbein & Knott, 2015, Agrawal et al., 2021). All in all, we see that, in practice, firms tend to waste scarce resources by persisting in pursuing bad ideas longer than they should.

I argue that an approach seeking to conserve resources with low upfront product investment can actually be a source of the very persistence this approach purports to avoid. In other words, the two key elements of the lean approach—minimal products and rapid pivoting—may not always work together as intended. Using a two-period model, I illustrate that when trying to sell a product with only

mediocre product investment, an entrepreneur is prone to be unable to learn whether a failure to sell is due to a lack of market demand (in which case she should pivot or quit) or insufficient investment in product development (in which case she should persist). I identify a “mediocrity trap” where an entrepreneur with a mediocre product can ex ante assign a higher expected value to persistence relative to pivoting and is, consequently, ex ante committed to persisting no matter what feedback she receives from the market. This suggests that persisting (rather than pivoting) in the face of negative feedback can be the rational path when following the generally recommended model of resource-light entrepreneurial learning. In other words, entrepreneurs persist when their products are too minimal to provide viable and actionable learning.

Consider the example of CellCo, a startup that sought to provide cellphone coverage to rural areas not served by large cell towers. The company quickly went to market after only investing enough to develop a product offering 2G coverage, aiming to learn and validate demand before investing in 3G or 4G technology. When network usage was low, it was unclear whether this negative market signal indicated a lack of market opportunity (e.g., no one wanted to use cell coverage in the area) or simply insufficient product quality (e.g., people did want to use cell coverage in the area, but 2G was not good enough). It was ambiguous. So the startup persisted in trying to validate the market by continuing to try to get customers to use its 2G coverage network, and ultimately shut down. I argue that product quality impacts what an entrepreneur is able to learn from market signals. Going to market quickly with a minimally developed product can actually impede learning and slow the resolution of uncertainty. I propose a strategy of balancing product development and sales over time to enable meaningful learning, especially learning from failure.

To empirically explore these theoretical insights, I construct a detailed panel of 1,445 enterprise software companies with API products. The empirical analysis is consistent with the model’s predictions: firms that take lower-quality products to market tend to persist in trying to sell even after failure. Further, I find that these firms are more likely to shut down in the long run and, contrary to

the supposed benefits of fast entry, actually take more time and expend more resources in the process. These results suggest that sufficient investment in product development is important to enable meaningful learning about whether product-market fit exists.

This study makes three main contributions. First, it contributes to the emerging theory-based view of entrepreneurship (Felin & Zenger, 2017, Wuebker et al., 2023) by suggesting that theoretical clarity about learning objectives and product quality thresholds can help entrepreneurs design more informative, critical market experiments, rather than relying solely on rapid iteration with minimal products (Valentine et al., 2024). Second, it contributes to the literature on entrepreneurial scaling (Knudsen et al., 2014, DeSantola & Gulati, 2017, Lee & Kim, 2024) by demonstrating how early product quality decisions can impact a venture's growth trajectory. Finally, it offers significant managerial implications and challenges prevailing managerial approaches to experiment and "fail fast" in a resource-efficient manner.

3.2 THEORETICAL BACKGROUND

3.2.1 PRODUCT AND MARKET AS KEY DIMENSIONS OF ENTREPRENEURIAL UNCERTAINTY

Uncertainty is broadly recognized as a fundamental component of entrepreneurship (March, 1991, Minniti & Bygrave, 2001, Politis, 2005). Strategy literature has long recognized two fundamental dimensions of uncertainty—product and market—which entrepreneurs must navigate (Helfat & Lieberman, 2002). While there are many different dimensions of uncertainty, these two dimensions are highlighted as fundamental, early-stage dimensions of uncertainty because entrepreneurs generally need to resolve product and market uncertainties before addressing other dimensions of uncertainty, such as ecosystem and institutional uncertainties (Moeen et al., 2020).

Navigating the product dimension of uncertainty encompasses developing the relevant techno-

logical capabilities to deliver a solution that performs the intended function (Agarwal et al., 2004). Navigating the market dimension of uncertainty encompasses identifying a group of customers with a specific need and associated willingness to pay (Adner & Zemsky, 2006). These two dimensions together constitute what is referred to as product-market fit: a working product that meets the needs of some market of real customers. Zellweger & Zenger (2023) define product-market fit formally as “the extent to which the opportunity belief held by the entrepreneur matches the market’s response” and Grégoire & Shepherd (2012) define this fit as the degree to which an opportunity’s “means of supply possesses qualities that meet the needs and requirements of a target market.” In short, product-market fit captures the essential interplay between resolving both the product and market dimensions of uncertainty—entrepreneurs must not only identify a genuine market need, but also develop a product solution which addresses that need.

The importance of both these dimensions is made more clear when considering ventures that succeed in navigating uncertainty along one dimension but fail in the other. For instance, Segway navigated product uncertainty to build what most would objectively evaluate as a high-quality transportation device, but failed to resolve the market uncertainty and find a viable market for this product (Kemper, 2003). Conversely, Theranos identified a legitimate market opportunity in compact blood testing, but failed to resolve the product uncertainty needed to develop a product that could accurately perform these tests (Carreyrou, 2018). These examples underscore that product and market represent distinct, separate sources of uncertainty, but ultimately both must be addressed for a business to be viable.

3.2.2 LEARNING APPROACHES FOR RESOLVING PRODUCT-MARKET UNCERTAINTY

Resource constraints present entrepreneurs with a fundamental challenge: they must resolve both product and market uncertainties with limited capital, time, etc. (Baker & Nelson, 2005, Clough et al., 2019). Although both dimensions need to be addressed to establish product-market fit, re-

source constraints force strategic choices about which uncertainties to prioritize and how to sequence learning efforts.

Many advocate for a market-focused approach, prioritizing allocating resources to early engagement with potential customers before investing in significant product development (Felin et al., 2020, Leatherbee & Katila, 2020). Building on von Hippel (1986)'s research on "lead users" and the concept of "market orientation" (Kohli & Jaworski, 1990) in the marketing literature, this approach encourages entrepreneurs to quickly develop initial product prototypes, gather customer feedback, and iterate rapidly. This perspective is most prominently embodied in practice in the Lean Startup methodology (Ries, 2011), which advocates for quick market testing of minimum viable products or MVPs (e.g., prototype products with only a subset of the envisioned features and functionality to conserve resources) and rapid pivots. Empirical support for this approach includes findings that entrepreneurs who engage with potential customers early can avoid over-investing limited resources developing products no one actually wants to buy (Camuffo et al., 2020).

However, there are also arguments for a different, product-focused approach, which emphasize investing in developing technological capabilities and more robust functionality before extensive market engagement (Utterback, 1974, Abernathy & Clark, 1985). This perspective questions whether customers can accurately evaluate or provide meaningful feedback on underdeveloped products, suggesting that most customers (unlike von Hippel's limited group of "lead users") assess products at face value rather than imagining what an early, bare-bones version of a product could potentially become with additional development. This approach suggests that the process of product development itself generates valuable information about potential product-market fit (Baldwin & Clark, 2006). It acknowledges that product quality often drives market acceptance (Agarwal & Bayus, 2002) and highlights how hard it can be to recover from reputational damage caused by launching underdeveloped products (Barnett & Pollock, 2012, Hsu, 2007).

3.2.3 PRODUCT QUALITY AND SIGNAL INTERPRETATION

THE SIGNAL CLARITY CHALLENGE

The dominant market-based approach, particularly in the Lean Startup methodology, rests on two fundamental prescriptions for finding product-market fit which are assumed to be complementary: (1) develop minimum viable products, and (2) learn rapidly from market feedback to quickly pivot away from bad ideas. These prescriptions aim to help entrepreneurs “fail fast” and avoid wasteful persistence. Recent research has begun to question the universal applicability of these principles (Felin et al., 2020). Valentine et al. (2024) suggest Lean methodologies may result in reactive pivots rather than strategic learning, and Agarwal et al. (2025) found Lean treatment groups perform worse than control groups. Building on these critiques, I propose that the two core elements of the Lean approach—minimal products and rapid pivoting—may not always be compatible.

This compatibility challenge stems from the relationship between product quality and signal noise. The concept of noisy signals in entrepreneurial learning is gaining attention (Posen et al., 2018), with Agrawal et al. (2021) highlighting how market signals convey information about both the idea and strategy. I extend this insight by arguing that product quality directly affects signal clarity and learning efficacy.

Consider an entrepreneur attempting to sell a new product. A successful sale provides an unambiguous positive signal: both product quality was sufficient and a market exists. However, a failure to sell yields different insights depending on product quality. If the entrepreneur had attempted to sell a hypothetically perfect product, failure would definitively indicate the absence of a market. In contrast, with a lower-quality product, failure to sell could be attributed to either lack of market demand or insufficient product development. Employing a Bayesian framework, bringing a lower-quality product to market decreases an entrepreneur’s ability to update her prior beliefs about whether product-market fit exists.

This asymmetry in signal clarity can be stated in the vernacular of established frameworks in strategy. Using [March \(1991\)](#)'s explore-exploit framework, this argument suggests that insufficient exploitation (product development) impairs effective exploration (market testing). Similarly, using absorptive capacity theory [Cohen & Levinthal \(1990\)](#), this argument suggests that investment in product development functions as a prerequisite for effective market learning.

THE MEDIOCRITY TRAP

I propose that the signal interpretation challenge creates what I call a “mediocrity trap”—a situation where entrepreneurs with underdeveloped products receive ambiguous negative market signals that paradoxically lead to persistence rather than pivoting.

When product quality is too low to generate clear market signals, entrepreneurs cannot effectively distinguish between product and market problems. This ambiguity can rationally lead to continued investment rather than pivoting, as entrepreneurs cannot rule out the possibility that market demand exists but their product is simply inadequate. Paradoxically, the very attempt to conserve resources through minimal product development—a central tenet of the Lean Startup approach—may result in prolonged resource expenditure through persistent market testing of inadequate products.

This insight directly challenges the presumed compatibility between minimal viable products and rapid pivoting. Rather than facilitating quick failure and efficient resource reallocation, minimally developed products may actually inhibit the learning necessary to make informed pivot decisions. In other words, the mediocrity trap suggests that two key components of a Lean approach—minimum viable products and “failing fast”—may not always be compatible. The formal model in the next section demonstrates how this mediocrity trap operates, showing that entrepreneurs who invest less in product development may not pivot, but rather persist in response to negative market signals.

3.3 SIMPLE MODEL

A simple model provides a structured analytical framework to explore when to try to sell a product to learn from customers versus when to continue developing the product further before trying to sell it.

The model considers an entrepreneur E operating over two periods with the objective of maximizing expected profit. In the first period, E must decide whether to go to market and learn by trying to sell her current product (“*Sell*”) or to instead invest in developing the product further (“*Develop*”). If she chooses *Sell* in the first period, she attempts to sell her product to some random customer and then, based on the outcome, updates her prior belief on whether the market exists and then uses this learning to decide whether to sell again (to a different random customer) or to pivot to another opportunity which has an alternative expected value a . Note that quitting and shutting down is a subcase of pivoting where the expected value in the second period is zero such that $a = 0$. In other words, E can choose to quit and with certainty incur no costs and no payoff in the second period. If she chooses *Develop* in the first period, she instead improves the product quality in the first period and then will try to sell her improved product in the second period.¹ So, in total, and as outlined in Figure 3.1, there are three potential paths over the two periods: (i) Sell, Sell, (ii) Sell, Pivot/Quit, and (iii) Develop, Sell.

Prior the start of the game, the existence of some real market opportunity m is determined exogenously. m is binary, equal to either 1 (if the market *does* exist) or 0 (if the market *does not* exist). m is unknown to the entrepreneur E . Instead, she has some previously formed prior belief μ as to whether the market exists. This belief $\mu \in [0, 1]$ represents the likelihood $m = 1$ or, in other words, how likely E believes it is that there is a market with viable customers who want to purchase her product. Also prior the start of the game, E also knows the starting quality of her product q . This quality $q \in [0, 1]$

¹We do not need to consider *Develop* in the second period. The reason is that given that this is the final period, it would never make sense to develop because E would never benefit from the investment. Also, because E has not received any feedback from the market after developing in the first period, she cannot update her priors and therefore would never choose to quit.

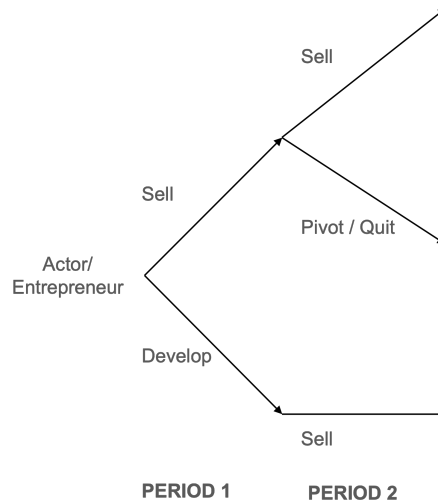


Figure 3.1: Model Overview. There are three possible paths: (i) Sell, Sell, (ii) Sell, Pivot/Quit, and (iii) Develop, Sell.

represents how well the product objectively performs a given task. I assume that quality is easy to observe and thus q is common knowledge and known to the entrepreneur E .

Given the separability of product and market, whether or not a customer chooses to buy a product (and the associated positive or negative market signal) is a function of both whether the market exists and the product quality. In other words, in order to sell two conditions must hold: there must be a market, and the product must be good enough to sell to that market. Since both conditions must hold, the overall probability of a successful sales effort is the product of the two. As such, the probability of a successful sales effort is mq and E 's expected likelihood of a successful sales effort is μq .

If E chooses to sell, while she has the chance to secure revenue, she also incurs some cost associated with selling (c_s), such as salespeople. This cost $c_s \in [0, 1]$ is normalized to the total possible revenue from a successful sale. If E chooses to develop, she incurs some cost associated with developing $c_d \in [0, 1]$ which is also normalized to the total possible revenue from a successful sale. However, if E chooses to develop, the quality of her product is improved by some Δq . Importantly, $\Delta q \leq 1 - q$ because overall quality is the probability of a sale conditional on a market existing and probabilities

must be smaller than 1.

3.3.1 ANALYSIS: EXPECTED VALUE OF EACH PATH

To decide whether to *Sell* or *Develop* in the first period, the entrepreneur E considers the expected value of each path: (i) Sell, Sell, (ii) Sell, Pivot/Quit, and (iii) Develop, Sell. Given that E is profit-maximizing, she will follow the path with the highest expected value and take the associated first period action (*Sell* or *Develop*).

LEARNING AND UPDATING PRIOR BELIEF

Learning is an important part of the expected value calculation and E 's subsequent actions. If E chooses *Sell*, she gets a signal from the market that can be used to update her prior belief μ as to the likelihood that a market for her product exists. Specifically, the signal she receives comes from whether or not she successfully sells her product.

A successful first-period sale allows E to update her prior belief μ to some updated belief μ_1 . Specifically, after getting the positive signal of a successful sale, she knows that $\mu_1 = 1$ because a successful sale can only happen when the market exists. In other words, $m q$ can only equal one if $m = 1$. If $m = 0$, then $m q$ will always equal zero and can never equal one. As a result, the expected likelihood of a successful sale in the second period after successfully selling in the first period is q (derived by plugging $\mu_1 = 1$ into E 's expected likelihood of a sale μq).

Upon failure to sell in the first period, E 's prior belief μ updates via Bayesian updating to $\mu_1 = \Pr(m = 1 | \text{fail}) = \frac{\Pr(\text{fail} | m=1) \Pr(m=1)}{\Pr(\text{fail})} = \frac{(1-q)\mu}{1-\mu q} = \frac{\mu - \mu q}{1 - \mu q}$. Looking at how E 's prior belief updates upon failure ($\mu_1 = \frac{\mu - \mu q}{1 - \mu q}$), we see that she will update her prior less for lower values of initial product quality q and more for higher values of q . A proof for Observation 1 can be found in Appendix Section C.1.1.

Observation 1: For lower values of initial product quality q , there will be a smaller difference between E 's prior belief μ and her updated belief after failure μ_1 .

For example, consider a product which is not developed at all such that $q = 0$. In this case $\mu_1 = \frac{\mu - \mu(0)}{1 - \mu(0)} = \mu$; if there is no prior investment in product development, a failure to sell provides no new information about the value of m and E cannot update her prior belief μ . In this case, there is no difference between μ and μ_1 . Alternatively, consider a product of perfect quality that performs its intended task perfectly such that no improvement is possible and $q = 1$. In this case $\mu_1 = \frac{\mu - \mu(1)}{1 - \mu(1)} = 0$; if the product is fully developed, a failure to sell allows the entrepreneur E to determine with certainty that $m = 0$ and a market does not exist.

For additional clarity, Figure 3.2 illustrates the information value of a negative signal for different values of initial product quality q . The x-axis shows different initial beliefs μ , and the y-axis shows the absolute change in E 's prior belief after getting a negative signal ($|\mu_1 - \mu|$). The different lines trace this difference in the absolute change in E 's belief for different values of q . This figure illustrates that lower values of q always lead E to change her prior by a smaller amount; there is less information value in the signal from trying to sell a lower-quality product. In other words, E learns less about whether there is a market (the signal is noisier) after failure when she tries to sell a lower-quality product because it is less clear whether the failure to sell is due to the lack of a market or low product quality.

As a result, the expected likelihood of a successful sale in the second period after a failed sales attempt in the first period is $\frac{\mu - \mu q}{1 - \mu q} q$ (derived by plugging in $\mu_1 = \frac{\mu - \mu q}{1 - \mu q}$ into E 's expected likelihood of a sale μq).

EXPECTED VALUE OF *SELL* PATH

The expected value of *Sell* in the first period is equal to the expected likelihood of a successful sale (μq) minus the cost associated with selling (c_s), such as salespeople. Altogether, the expected value of *Sell* in the first period is $\mu q - c_s$.




Figure 3.2: Updated Prior Belief μ_1 After Failure to Sell. The x-axis shows the entrepreneur E 's initial belief μ in the existence of a market for her product. The y-axis shows the absolute change in E 's belief after a failure to sell in the first period. The three lines represent three different values of initial product quality q . Note that smaller values of q always lead E to update her prior belief by a smaller amount.

The expected value of *Sell* in the second period depends on whether or not the first-period sale was successful. As noted previously, a successful first-period sale allows E to update her prior such that $\mu = 1$ because a successful sale can only happen when the market exists. As a result, the expected value of trying to sell in the second period after successfully selling in the first period is $q - c_s$.

Also as noted previously, upon failure to sell in the first period, E 's prior μ updates via Bayesian updating to $\mu_1 = \frac{\mu - \mu q}{1 - \mu q}$. Thus, the expected value of trying to sell in the second period after failing to sell in the first period is $\frac{\mu - \mu q}{1 - \mu q} q - c_s$. If this updated expected value of trying to sell is less than E 's alternative a , then E will not choose to sell again but will instead pivot to her alternative.

These outcomes inform the comprehensive expected value of choosing *Sell* in the first-period. Using backwards induction, I sum the first-period expected value, the second-period expected value after success weighted by the likelihood of success in the first period, and the second period expected value after failure weighted by the likelihood of failure in the first period:

$$\mathbb{E}(\text{Sell}) = \underbrace{\mu q - c_s}_{\text{Period 1 EV}} + \underbrace{\mu q(q - c_s)}_{\text{Period 2 EV After Success}} + \underbrace{(1 - \mu q) \max \left\{ \frac{\mu - \mu q}{1 - \mu q} q - c_s, a \right\}}_{\text{Period 2 EV After Failure}}.$$

EXPECTED VALUE OF *DEVELOP* PATH

The expected value of the *Develop* path is more straightforward, reflecting the net benefit of selling a product whose quality is improved by some amount Δq in the second period after absorbing the initial development cost c_d in the first period. If entrepreneur E chooses *Develop*, she forgoes the opportunity to sell and instead incurs some development cost $c_d \in [0, 1]$. Therefore, the first-period expected value is simply $-c_d$.

In the second period, E will then try to sell her product that is now improved by some amount Δq . The expected value of trying to sell the improved product in the second period is $\mu(q + \Delta q) - c_s$. Together the expected value of *Develop* can be written as follows:

$$\mathbb{E}(\text{Develop}) = \underbrace{-c_d}_{\text{Period 1 EV}} + \underbrace{\mu(q + \Delta q) - c_s}_{\text{Period 2 EV}}.$$

E will choose *Sell* or *Develop* based on which path has the highest expected value.

3.3.2 THE MEDIOCRITY TRAP: PERSISTENCE AND PRODUCT QUALITY

A key insight of this paper that follows from Observation 1 is that lower levels of initial product quality can make the signal from a failure to sell so noisy that E is unable to update her prior belief by a sufficient amount such that pivoting or quitting is the rational option. Instead, she will persist in selling again no matter what. This is captured in the following Proposition, with the proof outlined in Appendix Section C.1.2.

Proposition 1: *A lower product quality q may cause E —conditional on selling in the first period—to persist in the market and sell again after failure to sell (rather than pivot or quit after a failure to sell).*

Proposition 1 (Formalized): *For a non-empty part of the admissible (μ, c_s, a, q) parameter space, there exist \hat{q} with $0 < \hat{q} < 1$ such that for $q < \hat{q}$, $E(\text{Sell, Sell}) > E(\text{Sell, Pivot/Quit})$ and for $q > \hat{q}$,*

$$E(\text{Sell, Pivot/Quit}) > E(\text{Sell, Sell}).$$

I call this proposition the *mediocrity trap*: there exists a region of quality levels where E is trapped in selling again rather than learning and pivoting. Put in plain terms, if E tries to follow popular approaches and invests the minimum amount to develop an intermediate product of mediocre quality, rather than pivoting or quitting in response to negative market signals, her mediocre product may cause her to persist instead.

This mediocrity trap is illustrated in Figure 3.3 where I plot the impact of different levels of q on E 's expected value. In this specific case, I assume E 's initial belief $\mu = 0.5$, which means she believes it is equally likely that a market does and does not exist. I also assume $c_s = 0.1$, which would imply a 90% margin (this is a realistic, good margin for an enterprise software company). I set $a = 0$ such that E 's alternative is to quit and shut down.

Figure 3.3 shows that for high values of q , it is more attractive for E to quit after failure. However, for mediocre values of q , it is more attractive to persist and try to sell again after failure. Using specific numbers, when initial product quality is high at $q = 0.9$, $\mathbb{E}(\text{Sell, Pivot/Quit After Failure}) = 0.71$ which is greater than $\mathbb{E}(\text{Sell, Persist After Failure}) = 0.70$. Here, E will choose to quit after failure. In contrast, when initial product quality is mediocre at half the prior value such that $q = 0.45$, $\mathbb{E}(\text{Sell, Pivot/Quit After Failure}) = 0.20$ which is less than $\mathbb{E}(\text{Sell, Persist After Failure}) = 0.25$. Here, E will choose to persist after failure.

This is somewhat counterintuitive as one would typically expect companies with lower quality products to quit more quickly; instead, for the same level of confidence that a market exists and with the same sales costs, having a relatively lower-quality product can cause E to choose to persist rather than quit upon failure. In other words, when product quality dips below a certain point, then the signal from failure is not strong enough to allow E to update her prior belief to a sufficient degree such that she will quit following a failure to sell. Instead, she will choose to persist and continue to sell regardless of a negative market signal. This scenario contradicts the experimentation ethos of

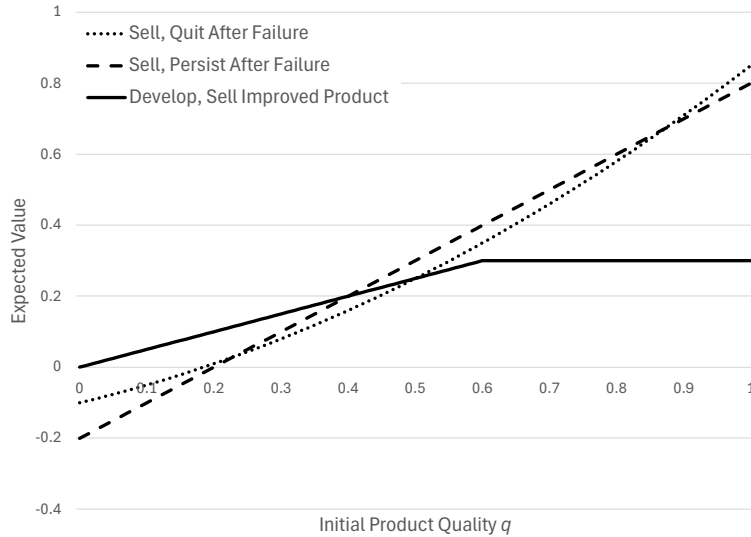


Figure 3.3: Model: Illustrative Case. The vertical axis *Expected Value* represents the ex ante expected value of choosing to sell versus choosing to develop in the first period determined through backwards induction. The horizontal axis *Initial Product Quality (q)* represents the quality of the product, which is known to the entrepreneur, at the start of the first period. In this illustrative case to enable plotting along these two dimensions, we assume the entrepreneur’s initial belief that there is a market $\mu = 0.5$, a cost of selling $c_s = 0.1$, a development cost $c_d = 0.1$, and an increase in quality from development $\Delta q = 0.4$, and that the alternative expected value is to quit with $a = 0$.

Note that the solid line represents developing in the first period. This line levels off after $q = 0.6$ because product quality cannot exceed 1; given that $\Delta q = 0.4$ and $q + \Delta q \leq 1$, the improved product quality will max out at 1 once q reaches 0.6. The dotted and dashed lines represent selling in the first period, but the rational entrepreneur will take the max of the two lines because she can choose whether to persist and sell again or quit in the second period. As such, the max of these two lines is what the entrepreneur will consider when comparing the expected value of selling versus developing in the first period.

This plot illustrates Proposition 1, showing that lowering the product quality can lead the entrepreneur to persist after failure rather than pivot/quit. It is initially counterintuitive that, holding the entrepreneur’s belief that there is a market, costs, etc. constant, lowering product quality would lead her to assume a higher expected value from persisting relative to quitting. Note that, in the counterfactual where the entrepreneur is forced to quit after failure, she will choose to develop more in the first period rather than sell.

failing fast to learn quickly, suggesting instead that products of mediocre quality may be what push entrepreneur E to market early due to the associated attractiveness of persisting despite failure.

Some comparative statics can be found in Appendix Section C.1.3. I show that the size of the range of values of initial product quality q for which E will fall into the mediocrity trap increases with m and decreases with c_s and a . In other words, E is more likely to fall into the mediocrity trap when she has a higher prior belief that the market exists, when cost of selling is low, and when the expected value of her alternative option from pivoting is low.

3.3.3 COUNTERFACTUAL COMPARISON

I now explore the counterfactual scenario where the entrepreneur E is forced to follow the “fast fail” logic and must always pivot after failure. As noted in Sections 3.2.2 and 3.2.3, the customer-centric, experimental approach is to invest less in product development and quickly try to sell an intermediate product so E can “fail fast” and avoid costly persistence. However, the model suggests that if E is forced to pivot or quit upon failure, she would actually choose to invest more often in product development before trying to sell; she would not try to sell a minimum mediocre product but would instead invest in developing a higher-quality product before trying to sell to the market.

Again referring to specific numbers, consider again the case referenced earlier where initial product quality is mediocre at $q = 0.45$, $\mathbb{E}(\text{Sell, Pivot/Quit After Failure}) = 0.20$, and $\mathbb{E}(\text{Sell, Persist After Failure}) = 0.25$. As illustrated in Figure 3.3, assume development cost $c_d = 0.1$ and the increase in quality from development $\Delta q = 0.4$. For the same $q = 0.45$, $\mathbb{E}(\text{Develop, Sell Improved Product}) = 0.23$. Here, E will choose to persist and sell again no matter what because it has the highest expected value. However, because $\mathbb{E}(\text{Develop, Sell Improved Product}) > \mathbb{E}(\text{Sell, Pivot/Quit After Failure})$, if E 's objective was instead to gain enough information to actually quit after failure, she would choose to develop the mediocre product further instead of going to market and trying to sell it now.

In summary, when E begins with a product of mediocre quality, the calculus of whether to *Sell*

or *Develop* can completely reverse if she were actually forced to pivot or quit after a failure to sell.²

3.4 HYPOTHESIS DEVELOPMENT

In this section, I develop hypotheses in the context of this paper’s empirical setting: enterprise software companies that offer Application Programming Interfaces (APIs) as their primary product. To provide clarity for readers who may be unfamiliar with this context, I will first explain some key concepts.

Enterprise software companies develop and sell software products and services to other businesses, rather than to individual consumers (Ethiraj et al., 2005, Pontikes, 2012). These products often aim to improve efficiency, productivity, or capabilities within the purchasing organization. APIs, or Application Programming Interfaces, are sets of protocols and tools that allow different software applications to communicate with each other (Zachariadis & Ozcan, 2017, Benzell et al., 2023). In essence, APIs act as intermediaries, enabling one piece of software to request and receive information or functionality from another.

For example, consider a university that uses a two-factor authentication system like Duo for secure login. Duo provides an API that allows the university’s IT systems to integrate Duo’s authentication service seamlessly. Another common example is Microsoft’s Active Directory, which many institutions use for single sign-on capabilities. Active Directory’s API allows various applications and services within an organization to verify user identities and permissions consistently. Importantly, these services are not developed by the schools themselves; instead, the schools use the APIs to access and integrate external technology into their own systems, leveraging specialized expertise and infrastructure provided by the API vendors.

In this context, acquiring a new customer is significantly different from the concept in consumer

²Modelt limitations and extensions are discussed in Appendix Section C.1.4

goods. When a consumer products company like a toothbrush manufacturer gains a new customer, it typically means a one-time purchase has occurred. In contrast, when an enterprise software company acquires a new customer, it often signifies the beginning of an ongoing relationship. These customers usually commit to subscription-based services or recurring purchases, with the expectation that they will continue to use and pay for the product over an extended period. This concept of a "recurring customer" is crucial, as it affects how these companies value new customer acquisitions and informs their sales and marketing strategies.

By hypothesizing in the context of enterprise software companies that offer API products, I aim to balance the generalizability of the findings with the need for a focused and controlled empirical setting. While the sample does not speak directly to all industries, the enterprise software sector represents a significant portion of the technology and venture capital landscape (with \$44B USD invested by VC funds in business-to-business software companies in 2023), making the findings relevant and valuable for a wide range of stakeholders. Furthermore, I will theorize why many of the relationships in this context are relevant across other industries.

3.4.1 MARKET SIGNAL: CUSTOMER ACQUISITION COSTS

In the enterprise software context, Customer Acquisition Cost (CAC) is defined as the total expenditure required to convert a potential client into a paying customer. CAC serves as a critical market signal and provides valuable insights into product-market fit. The underlying principle is that while sales can always be achieved given sufficient resource allocation, the efficiency of this expenditure, as reflected in the CAC, varies significantly based on several factors.

It is important to note that in the enterprise software industry, the cost of goods sold (COGS) for each incremental unit is typically quite low. The majority of costs are often associated with upfront fixed development expenses. However, a significant ongoing cost that can be directly tied to each sale is the cost of salespeople. In enterprise software, the sales process often involves complex negotiations,

demonstrations, and relationship-building, making salespeople a crucial and substantial part of the customer acquisition process.

A low CAC typically indicates strong product-market fit, suggesting that the offering resonates well with its target audience and requires minimal convincing or extended sales cycles. This efficiency often stems from a compelling value proposition, where the product's benefits are readily apparent and align closely with market needs. Conversely, a high CAC can signal problems with product-market fit. It may indicate underlying product issues or a misalignment between the offering and market needs, necessitating more extensive and resource-intensive efforts to educate and persuade potential customers. Put simply, it generally is not a question of whether a firm can sell a given product, but rather how much it will cost the firm to make each sale.

3.4.2 LEARNING FROM CAC AS MODERATED BY PRODUCT INVESTMENT

As discussed in the theory and model, what an entrepreneur does after getting a negative signal is key to unpacking what she learned from that signal. In this context, we can hypothesize on what an entrepreneur will do after selling and experiencing a high CAC (i.e., a negative signal that suggests there may be poor product-market fit) relative to how much she invested in the product previously. Learning (or a lack thereof) is then manifest as the entrepreneur either selling again or pivoting.

The theory predicts that entrepreneurs with lower-quality products may not update their beliefs in response to negative signals sufficiently to lead them to pivot or quit. Instead, they may persist in trying to sell again. This persistence manifests as continued investment in sales activities, such as hiring more salespeople, even in the face of high CAC. The rationale is that, with a mediocre product, the entrepreneur cannot definitively attribute high CAC to a lack of market demand versus insufficient product quality, leading the entrepreneur to continue their sales efforts rather than pivoting or quitting. In other words, and as suggested by the model, firms can be more likely to fall into the mediocrity trap and persist despite negative signals when trying to sell mediocre, low-quality products.

Hypothesis 6 *Enterprise software startups that get a negative signal of high CAC from selling a less-developed product are more likely to persist in selling again.*

Building on the logic of the mediocrity trap, we can extend the same reasoning to pivoting behavior. Pivoting involves moving away from one market and towards another, often in response to negative market feedback. However, the model suggests that entrepreneurs with less-developed products may be less likely to pivot, even when faced with negative signals such as high CAC. This is because the ambiguity in the market signal (high CAC) for a mediocre product makes it difficult for the entrepreneur to conclude definitively that the market does not exist, which would prompt her to pivot. Instead of recognizing the need to pivot, they may attribute the high CAC to factors that can be overcome with persistence. This reluctance to pivot is a direct consequence of the inability to learn effectively from the market signal due to the mediocre product quality.

Hypothesis 7 *Enterprise software startups that get a negative signal of high CAC from selling a less-developed product are less likely to pivot.*

Shutting down is a specific sub-case of pivoting (where the expected value is zero); the model suggests a similar relationship where E is less likely to shut down after receiving a negative signal from trying to sell a mediocre, low-quality product. The model suggests that this reluctance to shut down stems from the same mechanism that drives persistence and reduces pivoting: the inability to definitively interpret the high CAC as a signal of non-viability.

Hypothesis 8 *Enterprise software startups that get a negative signal of high CAC from selling a less-developed product are less likely to shut down.*

3.5 EMPIRICAL METHODS

3.5.1 CONTEXT

As noted previously, this study uses enterprise software companies that offer Application Programming Interfaces (APIs) as the empirical context. The enterprise software API context is particularly suitable for exploring the theory and hypotheses presented in this paper for several reasons. First, the relatively low technical risk in product development allows for a more direct examination of the relationship between product quality and market success. Unlike industries such as pharmaceuticals or advanced manufacturing, where product feasibility may be a significant barrier, enterprise software companies can generally build the products they envision. This means that the decision to try to sell a lower-quality product or to wait and develop a higher-quality offering is less confounded by technical constraints and more directly tied to strategic choice. Second, the success of enterprise software companies can be measured through various clear metrics, such as customer acquisition, retention, valuation, and key performance indicators specific to the enterprise software business model. These measurable market outcomes allow for testing the predictions in a way that might be more challenging in other contexts with less clear-cut measures of performance.

3.5.2 DATA AND SAMPLE

My primary dataset consists of Pitchbook venture data combined with LinkedIn employment data, BuiltWith API usage, and product architecture data. Pitchbook is an aggregator of data on private capital markets, including new ventures and venture capital investment in those ventures. Pitchbook uses a combination of technical software tools and a large human data operations team to assemble what is arguably the largest and most accurate database of startups and venture capital investment in those startups. LinkedIn contains data on the employees working at a given firm over time, as well

as the prior employment and educational experience of those employees. As outlined in [Stroube & Dushnitsky \(2023\)](#), BuiltWith data includes a comprehensive collection of internet technology data, profiling websites globally across multiple dimensions, including the technologies they employ, such as server infrastructure, analytics, frameworks, content management systems, and APIs. This dataset is generated by continuously scanning the web, capturing the technology adoption and usage patterns of millions of websites. The BuiltWith dataset offers a granular view of all the technologies websites use over time.

To define my sample I start with the 3,751 firms in the Pitchbook database who also have an API product captured on BuiltWith. To do so, I automate a search of each company in the Pitchbook database in the BuiltWith API database using each company's associated web domain. I then narrow my sample to the 1,445 firms with the "Information Technology" sector classification in Pitchbook given that these are the firms whose primary product is an API versus merely having an API as a secondary offering. For each firm, I collect the LinkedIn profiles of every individual who has ever worked at that firm, including both current and former employees. I then aggregate this data to construct various measures of employee composition at the firm-year level. I match the Pitchbook sample to the LinkedIn database using the firm webpage URLs. Next, I use the BuiltWith data to construct a panel at the website-year level, which captures all the websites (i.e., customers) using each API in each year. I then aggregate this data to the firm-year level to create measures of the customers each firm has in each year. Finally, I merge this firm-year panel with the previously constructed Pitchbook-LinkedIn panel. This approach allows me to create a comprehensive dataset that combines the Pitchbook venture data, LinkedIn employment data, and BuiltWith API usage and product architecture data at the firm-year level. The resulting panel dataset enables me to analyze the relationships between a firm's employee composition, customer base, product architecture, and performance over time.³

³All data is also available at the month level. However, I use year-level data given that sales cycles and other outcomes generally take multiple months and therefore a month-level analysis may not accurately capture performance.

3.5.3 VARIABLES

KEY DEPENDENT VARIABLES

The key insights outlined in my theoretical model and hypotheses center around what an entrepreneur decides to do after getting a negative market signal from trying to sell in the first period. As such, I construct measures to capture the second-period choices available to an entrepreneur after she tries to sell: *Sell Again*, *Pivot*, or *Quit*. I first construct a measure to capture a firm's persistence in selling again.

SELL AGAIN The nature of enterprise software makes salespeople a good proxy for a firm's strategic resource allocation to selling to customers. These products often involve more complex sales cycles that require personalized attention, tailored solutions, and significant customer education—all of which are effectively managed by salespeople. Unlike consumer products, enterprise software sales typically involve a more relationship-based approach, with longer decision timelines and negotiations that benefit from professional sales involvement. The decision to invest in a larger sales force indicates a strategic prioritization of sales efforts. The total number of salespeople in the subsequent year represents the firm's relative strategic prioritization and resources allocated to selling in that year. *Sell Again* is a leading variable equal to the total number of new salespeople hired by the firm in the subsequent year. As outlined in Appendix Section C.1.5, I find the results are robust to two alternative measures of persistence in selling again: total salespeople and salespeople attrition.

I then construct a set of measures to capture whether a firm pivots. A pivot involves both seeking new customer types and turning away from previously sought after customer types. To define customer types, I use the webpage data from BuiltWith that shows all the tools, technologies, and applications used to build each *customer's* webpage. How a customer's webpage is built is an especially relevant proxy for customer type when the customer is purchasing an API meant to fit in with the

ecosystem of other technologies on its website. By examining how a firm's customer type changes over time, I can measure what type of customer a firm is targeting and, when that changes, identify a pivot.

PIVOT: NEW CUSTOMER TYPES *Pivot: New Customer Types* is equal to the number of APIs used on new customers' webpages in the subsequent year that are not used by any customers in the current year, all relative to the total number of current customer APIs: in other words, how distinct the new types of customer webpages are relative to the existing customer base.

PIVOT: CHURNED CUSTOMER TYPES *Pivot: Churned Customer Types* is equal to the number of APIs used on current customers' webpages that are not used by any of a firm's customers in the subsequent year, all relative to the total number of current customer APIs: in other words, the degree to which a firm is turning away from a certain customer type.

I also construct a measure to capture whether a firm shuts down.

SHUTDOWN *Shutdown* is a binary indicator of whether the firm shuts down in a given year, equal to 1 in the year the firm's total employees drop below 50% from the peak.⁴ I argue that this approach to determining firm shutdown, enabled by the unique dataset on employees, is an improvement over the approach used by many other studies that rely on reporting of firm shutdown despite firm shutdown being much less likely to be reported relative to funding rounds, acquisitions, IPOs, etc. For example, PitchBook reports that only nine of the firms in the sample shut down during my period of interest. However, this is inconsistent with the well-established reality of the frequency of startups shutting down. Using my new measure of firm shutdown based on customer data, I find that 91 firms shut down, which is much closer to what one might expect in the first five years.⁵

⁴I also use different cutoffs in supplemental analyses to test for robustness.

⁵I find that the results hold if only using shutdowns reported in Pitchbook, though given that they are much less frequent the effect sizes and significance are smaller.

MAIN INDEPENDENT VARIABLES

CUSTOMER ACQUISITION COST (CAC) Generally, CAC is calculated by taking all sales expenses during a given period and dividing by the total number of new customers acquired during the same period. In the enterprise software context, the sales expense is generally just the sales employees. Therefore, I proxy *CAC* as the total number of salespeople employed by a given firm in a given year divided by the total new customers acquired by a given firm in a given year. A higher CAC is a negative signal because it means the firm must expend a larger amount of resources to acquire each new customer.

In calculating CAC, I employ the logarithmic division rule to address potential computational issues. Traditionally, CAC would be calculated as Total Salespeople Months (total months of all sales employees at a given firm in a given year) divided by Total New Customers (total new customers acquired during the year). However, this approach becomes problematic when New Customers equals zero, resulting in an undefined value. To circumvent this issue, I leveraged the property of logarithms where the log of a quotient equals the difference of the logs: $\ln(a/b) = \ln(a) - \ln(b)$. Specifically, I define *Customer Acquisition Cost (CAC)* as $\ln(\text{Total Salespeople Months} + 1) - \ln(\text{New Customers} + 1)$. Adding 1 to each term before taking the logarithm ensures that the argument is always positive, even when the original value is zero. This transformation normalizes the distribution of the CAC variable and resolves the division-by-zero problem, making it more suitable for statistical analyses. This approach allows for a more robust calculation of CAC across all firms in the dataset, including those that may not have acquired any new customers in a given period.

PRODUCT DEVELOPMENT The *Product Development* variable, measured as the cumulative sum of product developer months in all prior years, serves as a proxy for the firm's commitment to product development and investment in product quality. The underlying assumption is that firms investing more resources in their product development teams are likely to create higher-quality products. This

is similar to the approach used by many papers that use input measure as a proxy for quality or performance: research and development (R&D) expenditure is often used as a proxy for firm innovation (e.g., Rothaermel & Hess, 2007), and years of education and/or experience are often used as proxies for human capital quality (e.g., Hitt et al., 2001).⁶

There are potential concerns around a set of plausible confounding factors relative to *CAC* and *Product Development*, including number of products, competition, price, etc. In Section 3.5.4, I outline various robustness tests and theoretical explanations for why *CAC* is a valid measure.

CONTROL VARIABLES, FIXED EFFECTS, AND STATISTICAL MODEL

To account for the potential difference in impact of firms with more or less money available over time, and therefore different degrees of resource constraint over time, *Funding* is equal to the total cumulative amount of money raised by a firm at a given point in time measured in millions of USD and is included in all regressions.

To account for the changing size of the firm over time and associated changes in workforce capacity, organizational structure and complexity, etc., *Total Employees* is equal to the total months of employee labor used by a firm in a given year and is included in all regressions.

In addition, I also include a set of fixed effects to control for potential unobserved heterogeneity correlated with the variables of interest. *Firm* fixed effects allow for the comparison of within-firm variation in value over time by controlling for all unobserved heterogeneity at the firm level, including variation in industry, novelty, founder characteristics, etc. *Year* fixed effects account for the concern that certain outcomes may be more likely to occur earlier or later in a firm's life-cycle (e.g., a firm is more or less likely to hire salespeople right after founding versus many years after founding).

⁶These measures of course have limitations, as the investment in inputs does not necessarily equate to the quality of their output. To try to address this concern, Appendix Section X weights each product employee by their cumulative years of prior product development experience across prior jobs.

Given that this data is at the firm-year level, in order to account for potential correlation in the error term across firm age year, I generally cluster standard errors at the year level.

Also, in regression analyses I take natural logs of most variables (aside from binary indicators), as they approximately follow a log-normal distribution. These log variables are labeled with an Ln preceding the variable name in regression tables.

SUMMARY STATISTICS

As outlined in Table 3.1, the dataset comprises 1,445 firms observed over 9,343 months starting two years before the founding date to account for pre-founding activity and continuing through five years post-founding (with an average firm age of 1.9 years post-founding) in order to capture the key initial years of firm activity while eliminating the large variance and outliers late in firms' life-cycles. The summary statistics provide insights into various aspects of these firms' operations and performance.

Table 3.1: Summary Statistics. 1,445 Firms with observations over 9,343 months.

Statistic	Mean	St. Dev.	Min	Max
Sell Again	2.73	7.78	0	259
Pivot: Churned Customer Types	0.05	0.10	0	1
Pivot: New Customer Types	1.51	19.32	0	1,087
Shutdown	0.01	0.10	0	1
Product Development	62.92	285.72	0	10,838
CAC	0.66	2.33	-10.74	7.61
Funding	10.75	202.94	0	18,254.11
Total Employee Months	155.59	503.35	0	18,187
Year	1.88	1.97	-2	5

On average, firms hire 2.73 new salespeople in a given year. The average firm employs about four salespeople (47.8 salespeople months divided by 12 months in a year). However, this has a very large standard deviation of ~ 12 salespeople. The average firm has ~ 13 employees (155.6 employee months divided by 12 months), but again with a very large standard deviation of ~ 42 employees.

Pivot measures reveal that firms are more likely to explore new customer types (mean = 1.51) than to churn existing ones (mean = 0.05). However, the high standard deviation for new customer types (19.32) suggests significant variability in this behavior across firms. *Shutdown* has a low mean (0.01). However, it should be noted that this value only occurs once in the lifetime of each firm that fails in the panel. In practice, 31.3% of firms are marked as shutting down at some point.

Table 3.2 shows the correlation matrix between all numeric variables. *CAC* has a weak to moderate positive correlation with *Sell Again: New Salespeople* (0.134) and with *Sell Again Alt Measure: Total Salespeople Months* (0.155), which aligns with the idea that larger sales forces might lead to higher acquisition costs. *Total Employee Months* is strongly correlated with *Product Development* (0.730), consistent with the idea that larger firms have invested more in product development and have larger salesforces. Interestingly, the pivot measures show weak correlations with most other variables, suggesting that pivoting behavior might be influenced by factors not captured in these primary metrics. In addition, the measures of firm shutdown show very weak correlations with most variables, which might be due to its low frequency in the panel.

Table 3.2: Correlation Table. Pairwise correlations.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
(1) Sell Again	1								
(2) Pivot: Churned Customer Types	0.054	1							
(3) Pivot: New Customer Types	0.002	-0.016	1						
(4) Shutdown: Employees	-0.030	0.135	-0.002	1					
(5) Product Development	0.352	0.091	0.003	-0.007	1				
(6) CAC	0.134	-0.296	0.032	-0.119	0.105	1			
(7) Funding	0.087	0.013	0.001	0.006	0.158	-0.036	1		
(8) Total Employee Months	0.538	0.089	0.001	-0.019	0.730	0.096	0.219	1	
(9) Year	0.203	0.410	0.045	0.123	0.154	-0.098	0.058	0.192	1

3.5.4 EMPIRICAL LIMITATIONS, ROBUSTNESS TESTS, AND FUTURE DIRECTIONS

While the empirical analysis provides evidence consistent with the theoretical model, several important limitations should be acknowledged. This section discusses these limitations, the robustness tests conducted to try to address them, and some opportunities for future research.

MEASUREMENT AND CONSTRUCT ALIGNMENT

First, the empirical analysis uses proxies for the theoretical constructs and the distinctions between theory and empirics should be noted. For example, the measure of persistence—hiring additional salespeople—represents an intensification of sales efforts rather than merely continuing existing activities as conceptualized in the model. This stronger operationalization provides a clearer signal of strategic commitment and is more readily observable in the data. To validate this choice, alternative measures of persistence including total salespeople months and salespeople attrition are also tested (Appendix Section C.1.5), with consistent results across specifications.

Similarly, CAC is continuous versus the model's simplified binary feedback structure. Several robustness checks help validate CAC as a meaningful market signal, addressing concerns about other customer acquisition channels (Appendix Section C.1.9), changes in price (Appendix Section C.1.12), and the possibility that high CAC represents strategic investment rather than poor product-market fit (Appendix Section C.1.13). These tests consistently support the interpretation of CAC as a valid market signal.

ENDOGENEITY AND IDENTIFICATION CHALLENGES

Both product development investments and sales resource allocation represent endogenous strategic choices potentially influenced by unobserved factors. While firm fixed effects control for time-invariant heterogeneity, they cannot address time-varying confounders or reverse causality. To par-

tially mitigate these concerns, we investigated the influence of changes in firm strategy (Appendix Section C.1.10) and accounted for the impact of competition using two different clustering techniques (Appendix Section C.1.7), with the main findings remaining robust.

The use of firm fixed effects restricts analysis to within-firm variation over time rather than between-firm differences. This approach differs from the model's conceptualization of an entrepreneur making an initial product quality decision, then observing market signals and deciding whether to pivot. However, this panel approach offers significant analytical advantages by controlling for unobserved firm-level heterogeneity. The fact that we find results consistent with the theoretical predictions despite using this different analytical approach suggests that the effect appears to operate within firms over time.

The appendix further addresses potential selection concerns related to data sources (Appendix Section C.1.6), with results holding when restricting the sample to post-2012 observations. I also investigate how customer size (Appendix Section C.1.11) and the presence of multiple products (Appendix Section C.1.8) might affect our results.

FUTURE RESEARCH DIRECTIONS

Despite these robustness tests, several limitations remain that offer promising directions for future research. Experimental or quasi-experimental approaches could provide stronger evidence. For instance, examining exogenous shocks to product development capabilities or conducting field experiments that randomly vary product investment before market testing would address the endogeneity concerns inherent in observational studies.

The current study focuses exclusively on enterprise software companies offering API products, limiting generalizability to other contexts. Comparative analyses across different industries with varying technological characteristics, development timelines, and feedback mechanisms would help identify boundary conditions and refine our understanding of when product quality most critically influ-

ences entrepreneurial learning.

The long-term effects of early resource allocation decisions (Appendix Section C.1.14) suggest that initial emphasis on product development versus sales may have significant performance implications. Future research could more systematically track how entrepreneurs learn to escape this trap over time, potentially developing interventions that help entrepreneurs more effectively balance product development with market testing.

In conclusion, while the empirical analysis provides evidence consistent with the model and theory, these findings should be interpreted with appropriate caution given the methodological limitations outlined above. Nevertheless, the consistency of results across multiple robustness tests strengthens confidence in the core finding: entrepreneurs with less-developed products tend to persist despite negative market signals.

3.6 RESULTS

Hypothesis 6 predicts that, upon receiving a negative signal of high CAC, firms without meaningful investment in product development will be more likely to persist and sell again in spite of the negative signal. The regression analysis in Table 3.3 shows support for this hypothesis.

The regression analysis reveals a positive relationship between *CAC* and *New Salespeople* ($\beta = 0.018, p = 0.144$). This suggests that a 100% increase in CAC is associated with an approximate 1.3% increase in the number of new salespeople hired in the next period.⁷ This indicates that firms tend to expand their sales force in response to higher customer acquisition costs, consistent with persistence in selling efforts.

The interaction term between *Product Development* and *CAC* is negative ($\beta = -0.007, p = 0.071$). This suggests that higher product investment mitigates the impact of CAC on hiring new

⁷Effect size of log-log transformed variable pair is calculated as $2^{0.018} - 1 \approx 0.0126$ here and in a similar manner subsequently.

Table 3.3: Tendency to Persist and Sell Again. Ordinary least squares (OLS) estimation. Standard errors (in parentheses) are clustered at the year level. p -values are shown in brackets.

	Sell Again		
Ln Product Development	0.185 (0.004) [0.000]	0.092 (0.019) [0.002]	0.050 (0.012) [0.004]
Ln CAC	0.089 (0.006) [0.000]	0.033 (0.014) [0.050]	0.018 (0.011) [0.144]
Ln Product Development \times Ln CAC	-0.006 (0.002) [0.000]	-0.008 (0.003) [0.041]	-0.007 (0.003) [0.071]
Ln Funding			0.095 (0.030) [0.015]
Ln Total Employees			0.110 (0.018) [0.000]
(Intercept)	0.317 (0.012) [0.000]		
Observations	9343	9343	9343
R^2	0.239	0.617	0.627
FE: Year		X	X
FE: Firm		X	X

salespeople. To fully offset the effect of a 100% increase in CAC on new salespeople hires, product investment would need to increase by approximately 185.4%.⁸ This indicates that while higher CAC generally encourages startups to hire more salespeople, product investment can offset this effect.

Table 3.4: Tendency to Pivot or Shutdown. Ordinary least squares (OLS) estimation. Standard errors (in parentheses) are clustered at the year level. p -values are shown in brackets.

	Pivot: Churned Customer Types	Pivot: New Customer Types	Shutdown
Ln Product Development	-0.001 (0.001) [0.526]	0.001 (0.018) [0.944]	-0.008 (0.003) [0.019]
Ln CAC	-0.014 (0.000) [0.000]	0.030 (0.018) [0.150]	-0.010 (0.003) [0.014]
Ln Product Development × Ln CAC	0.001 (0.000) [0.063]	0.015 (0.004) [0.005]	0.002 (0.001) [0.009]
Ln Funding	0.002 (0.001) [0.021]	0.020 (0.010) [0.075]	0.007 (0.003) [0.040]
Ln Total Employees	0.008 (0.001) [0.000]	-0.024 (0.037) [0.533]	-0.038 (0.018) [0.072]
Observations	8690	7929	9608
R^2	0.586	0.371	0.291
FE: Year	X	X	X
FE: Firm	X	X	X

Hypothesis 7 predicts that enterprise software startups receiving a negative signal of high CAC from selling a less-developed product are less likely to pivot. This hypothesis is supported by the results in Table 3.4.

The analysis shows a positive relationship between *CAC* and *Pivot: New Customer Types* ($\beta = 0.030$, $p = 0.150$). A doubling of CAC is associated with a 2.1% increase in the measure of new customer types. The interaction term between *Product Development* and *CAC* is positive ($\beta = 0.015$, $p = 0.005$), indicating that higher product development investment meaningfully amplifies the effect

⁸ Calculated as $2^{(0.018/0.007)} - 1 \approx 1.8544$.

of CAC on new customer types. Firms with more investment in product development seem to pivot and explore more new types of customers in the face of failure.

For *Pivot: Churned Customer Types*, there is a negative relationship with CAC ($\beta = -0.014$, $p < 0.01$). A doubling of CAC is associated with a 0.97% decrease in the measure of pivoting away from existing customer types. This suggests that firms are less likely to turn away from existing customer segments when faced with high acquisition costs for their current product. The interaction term between *Product Development* and CAC for *Pivot: New Customer Types* is positive ($\beta = 0.001$, $p = 0.063$), indicating that higher product development investment can mitigate the negative effect of CAC on pivoting away from existing customer types.

These results are consistent with the proposed mediocrity trap, where firms with less-developed products may persist in the face of negative signals rather than pivoting or exiting.

3.6.1 POST HOC ANALYSIS: LONG-TERM IMPLICATIONS OF EARLY SALES AND PRODUCT INVESTMENT

While the primary analysis focuses on the next-year effects of product quality on entrepreneurial learning and decision-making, I also conduct a post hoc analysis to explore potential long-term consequences of early prioritization of trying to sell versus investing in product development. This analysis aims to understand how firms' early resource-allocation decisions and associated strategic priorities may impact their future performance. In Table 3.5, I look at the long-term impact of a firm's total sales employees (as a proxy for resources allocated to trying to sell) versus product development employees (as a proxy for resources allocated to product development) during its first two years.⁹ Specifically, I show that early prioritization of sales alone leads to fewer total customers and a decreased likelihood of becoming a unicorn in the long run; only when balanced with sufficient product development are sales efforts in the first two years associated with more total customers and a higher likelihood of

⁹I show that the results also hold if calculated based on first-year investments.

becoming a unicorn. Similarly, I find evidence that sales must be balanced with ample product development in order to decrease the likelihood of shutting down and, conditional on needing to shut down, not waste resources by delaying exit but rather truly fail fast. Appendix Section C.1.15 explores these analyses in greater detail.

Table 3.5: Initial Sales and Product Investment and Long-Term Outcomes. Ordinary least squares (OLS) estimation. Standard errors are clustered at the industry level and shown in parentheses. *p*-values are shown in brackets.

<i>Total Values in First Two Years</i>	Ln Total Customers	Unicorn	Shutdown	Years to Shutdown
Ln Salespeople	-0.043 (0.018) [0.076]	-0.014 (0.001) [0.002]	0.000 (0.000) [0.421]	0.011 (0.044) [0.810]
Ln Product Developers	-0.161 (0.002) [0.000]	-0.011 (0.001) [0.004]	0.009 (0.000) [0.000]	-0.058 (0.034) [0.159]
Ln Salespeople × Ln Product Developers	0.022 (0.002) [0.000]	0.004 (0.000) [0.003]	-0.001 (0.000) [0.000]	-0.024 (0.010) [0.063]
Ln Total Employees	0.232 (0.025) [0.000]	0.067 (0.001) [0.000]	-0.018 (0.002) [0.000]	0.411 (0.049) [0.001]
Ln Funding	0.112 (0.034) [0.030]	0.001 (0.000) [0.032]	-0.001 (0.001) [0.140]	-0.376 (0.087) [0.012]
Observations	1445	850	1445	305
<i>R</i> ²	0.117	0.168	0.068	0.747
FE: Industry	X	X	X	X
FE: Founding Year	X	X	X	X
FE: Country	X	X	X	X
FE: Technical Founder	X	X	X	X

3.7 DISCUSSION

3.7.1 THE THEORY-BASED VIEW

This study contributes to the theory-based view of entrepreneurship (Felin & Zenger, 2017, Wuebker et al., 2023), which conceptualizes entrepreneurs as scientists who should “thoughtfully craft...a theory amenable to efficient testing and careful feedback interpretation” (Zellweger & Zenger, 2023, p.

398). This view emphasizes the importance of entrepreneurs developing causal theories about paths to value creation. My findings suggest that the quality of the testing apparatus itself—the product—fundamentally impacts an entrepreneur’s ability to test theories and learn from market experiments. The mediocrity trap demonstrates how inadequate product development can undermine the theory-testing process, creating conditions where the entrepreneur will persist regardless of an experiment’s outcome.

Both this study and the theory-based view have clear ties to Popperian falsifiability (Popper, 1969, p. 39). Popper’s principle that scientific theories must be testable and potentially refutable is embraced throughout the theory-based view’s conception of entrepreneurship as a scientific process. My findings enrich this perspective by revealing how product quality serves as a crucial moderator of falsifiability in entrepreneurial contexts. Mediocre products can generate ambiguous signals that limit an entrepreneur’s ability to determine whether negative feedback falsifies her market theory or merely reflects inadequate product implementation. By identifying product quality as a consideration in effective falsification, this study supports and extends the theory-based view’s application of Popperian principles to entrepreneurship.

In addition, both this study and the theory-based view allude to the concept of *experimentum crucis*—or “critical experiment”—used widely by scientists since Francis Bacon and Isaac Newton to decisively determine the validity of a theory. Zellweger & Zenger (2023) describe how entrepreneurial beliefs direct attention to “what has to be critically true for the product to ‘work,’” highlighting the importance of testing these critical assumptions (p.385). A critical experiment enables researchers to distinguish between competing theories through tests that provide unambiguous results. I propose that for entrepreneurial market experiments to function as true critical tests, they require products of sufficient quality to generate clear, actionable results.

These insights join a growing body of work within the theory-based view that explores boundary conditions of the lean startup. The mediocrity trap concept suggests that entrepreneurial theory test-

ing requires products of sufficient quality to generate decisive evidence—a finding that connects to recent work on pivot quality (Valentine et al., 2024) and startup performance (Agarwal et al., 2025). By demonstrating how product quality enables or constrains entrepreneurial learning, this study enriches the theory-based view and suggests promising research directions examining the interplay between product quality thresholds and experimentation effectiveness. Future work might explore how entrepreneurs can determine the minimum level of product quality needed for their critical experiments and whether this threshold varies across different industry contexts or types of entrepreneurial theories.

3.7.2 SCALING ENTREPRENEURIAL VENTURES

Despite the critical role of scaling in the entrepreneurial process, academic research on this topic is historically less explored in the strategic management literature (Knudsen et al., 2014). DeSantola & Gulati (2017) conceptualize scaling as the process of establishing and expanding an organization to serve a growing customer base. They highlight the organizational challenges and key inflection points that startups face during the scaling process. More recently, Lee & Kim (2024) provide large-scale empirical evidence on the timing of scaling, finding that startups that scale too early, as indicated by their first job postings for managers and sales personnel, are more likely to fail. They attribute this to the tension between reducing imitation risk and increasing commitment risk.

The present study extends this discussion by focusing on the specific balance between product quality and the timing of pushing sales, a crucial component of scaling. The essence of the mediocrity trap is that lower-quality products may push entrepreneurs to go to market and sell despite not being able to learn from the market. In other words, this study suggests that having a lower-quality product may be part of what pushes firms to scale prematurely. This insight underscores the importance of achieving a sufficient level of product quality before expanding sales. This is one specific mechanism that supports the managerial recommendation to “nail it, then scale it” (Furr & Ahlstrom, 2011) Fur-

thermore, this study contributes to the broader understanding of the complexities and challenges associated with scaling in entrepreneurial ventures (Hambrick & Crozier, 1985, Nicholls-Nixon, 2005).

3.7.3 MANAGERIAL IMPLICATIONS

This study challenges conventional entrepreneurial wisdom by revealing how mediocre, less developed products can paradoxically lead to persistence rather than pivoting. For entrepreneurs and investors, the central implication is the need to carefully balance product development and market testing, rather than rushing minimally viable products to market. To use managerial vernacular, it is important to think carefully about what “viable” means in the concept of a minimum viable product—while making something “minimum” is relatively straightforward, this research suggests that careful attention should be given to whether a product is not only minimum, but also viable. This study suggests that viable products provide actionable learning. Entrepreneurs and investors should think carefully about whether a failure to sell a prospective MVP would actually allow learning about the market. Will it provide a clear answer to a critical assumption, or will it leave room for ambiguity? A truly viable product must be developed enough to yield clear market signals that can definitively validate or invalidate key assumptions. Entrepreneurs and investors should establish and meet specific criteria for viability before market testing.

This principle extends beyond entrepreneurship to other domains where experimentation and learning are critical. In academic research, scholars face similar trade-offs—scholars should think carefully about what constitutes a minimum viable idea and the extent to which to develop an idea before sharing it to get feedback. In creative fields like film and television, producers must decide whether early cuts or scripts are sufficiently developed to generate meaningful audience reactions. In corporate innovation, product managers must determine when prototypes are sufficiently developed to yield valid customer insights. In each context, insufficient development can obscure learning; effective learning requires investments in quality that enable true critical tests.



Entrepreneurial Learning and Strategic Foresight

A.1 ONLINE APPENDIX

A.1.1 QUALITATIVE INTERVIEWS

We conducted a program of qualitative interviews after deductively theorizing hypotheses and testing

those hypotheses empirically to determine the mean effect. These interviews were not used to draw inductive conclusions. Instead, the purpose of these interviews was to better understand and confirm the mean effect observed in the quantitative data. Given that intent, we identified and reached out to entrepreneurs in our sample who match the observed empirical trends to gain additional color on what drove those outcomes. These interviews confirmed the mechanisms outlined in our theory, which we established prior to conducting the interviews.

We reached out to entrepreneurs via LinkedIn messages, company contact forms, Kickstarter messages, or email, depending on what was available in each case. We received 18 responses to our outreach. Two of these entrepreneurs were not willing to participate and five dropped out in the scheduling process. In the end, we conducted interviews with 11 entrepreneurs from our sample. Table C.1 and Table C.10 provide summary information and statistics for each of these entrepreneurs. We conducted all interviews via Zoom videoconferencing except for one interview conducted via Google Hangouts per the entrepreneur's request. All interviews conducted via Zoom were recorded after obtaining the entrepreneur's verbal consent and transcribed for review. Interviews were conducted between January 29, 2020 and June 24, 2020. The interviews were scheduled for 30 minutes, with most interviews lasting between 30 and 45 minutes.

Each interview followed a semi-structured format. Interview questions covered a broad range of topics roughly mirroring the topics outlined in our theory, hypothesis development, and supplemental analysis, as well as general background. We asked entrepreneurs: why they decided to launch crowdfunding projects on Kickstarter (and why they stopped); what they learned implementing an earlier project and how that impacted future projects; what their process was for determining the predicted product specification and predicted timeline; what unforeseen challenges they encountered; what they did when they encountered unforeseen challenges, etc.

Table A.1: Qualitative Interview Sample. Indexed set of interviewed entrepreneurs. All interviews lasted 30 to 45 minutes and were conducted between January 29, 2020 and June 24, 2020.

	Entrepreneur	Product Subtype	Location	Location Type
1	3DPrintCo CEO	3D Printing	Spicer, MN	Town
2	AccessoryCo CEO	Hardware	Izhevsk, Russia	Town
3	CircuitsCo CEO	DIY Electronics	North Sydney, AU	Suburb
4	ElectronicsCo CEO	DIY Electronics	Preston, UK	Suburb
5	GPSCo CEO	Gadgets	Seattle, WA	Town
6	LEGOCo CEO	Hardware	London, UK	Town
7	MaskCo CEO	Wearables	Montreal, Canada	Town
8	MusicCo CEO	Sound	Austin, TX	Town
9	SecureCo CEO	Hardware	Dublin, Ireland	Town
10	TabletCo CEO	Hardware	Beijing, China	Town
11	WidgetsCo CEO	DIY Electronics	Boulder, CO	Town

Table A.2: Qualitative Interview Sample Summary Statistics. Indexed set of interviewed entrepreneurs. All values are averaged across all the projects in the sample by the individual entrepreneur. *Funding Threshold*, *Funding Exceeded*, and *Funding Backers* are all in thousands.

	Funding Threshold	Funding Exceeded	Funding Backers	Unforeseen Interdependencies	Delay Duration	Predicted Time	Actual Time
1	12.5	22	0.1	5.0	33.6	96.8	130.4
2	230.0	504	3.9	14	381.9	148.3	530.2
3	1.1	12	0.5	3.5	27.8	89.5	117.3
4	3.1	14	0.5	1.0	3.7	51.8	55.5
5	10.0	145	2.3	3.0	37.0	59.0	96.0
6	76.4	128	1.1	4.5	154.5	79.6	234.1
7	8.6	165	3.3	1.0	53.5	91.7	145.2
8	42.5	108	0.9	3.5	128.7	107.4	236.1
9	86.7	121	0.5	4.0	359.5	165.6	525.1
10	75.0	929	11.6	2.0	14.0	51.9	65.9
11	8.2	25	0.3	0.5	4.9	112.3	117.2

A.1.2 COMPLEXITY IN CROWDFUNDED HARDWARE TECHNOLOGY PROJECTS

We provide contextual evidence that: (1) entrepreneurs tend to add multiple features to subsequent projects, and (2) even relatively “low” complexity projects in our sample were still highly complex and required addressing a large number of interdependencies.

WITHIN ENTREPRENEUR: TENDENCY TO ADD MULTIPLE FEATURES

Close examination of specific products in our sample make it clear that entrepreneurs generally add multiple features as they advance from project to project, leading to increased complexity across an entrepreneur’s hardware technology crowdfunding projects. Continuing from Figure 2.2, we use the example of Plum Geek Robotics to illustrate this progression across various quantitative and qualitative measures. Figure A.1 presents a detailed breakdown of the product features in Plum Geek Robotics’s first and third projects.

The first project, Ringo Robot, has six colored lights, multi-frequency sound abilities, a light sensor, 360-degree visibility edge sensors, an accelerometer, a gyroscope, etc. By the time the company reached their third project, the Spirit Rover, the product had significantly more features, with a total of 27 colored lights, new computer vision capabilities, retractable gripper arms, wireless network capabilities, etc.

If we consider this example in the language of the NK model (Kauffman, 1995), N increased rapidly from the first to the third project. Based on our oversimplification of the product, we might imagine that N went from 5 to 15 to 35 across three projects. This increase in N is one way that the total number of interdependencies can increase rapidly from project to project.

ACROSS ENTREPRENEUR: “HIGH” COMPLEXITY ON DIFFERENT PROJECT TYPES

The hardware technology projects in our sample generally have high complexity, with both a large number of features and high level of interdependencies. That said, there is certainly variation in both the number of features and the level of interdependencies among those features—particularly when the features are modularized and separate from one another.

Before we highlight our context, we want to start with an extreme example of a type of product that is outside of our sample: board games. Board games are a popular crowdfunded project type, and it would certainly be the case that the printing of a game board is quite modular and separate from the die casting of game tokens. In this example, the entrepreneur would face limited interdependence when executing on the game board and the game tokens, except perhaps when packaging them together in a box, e.g., if the game tokens might scratch the game board in shipping. This example illustrates a product where both components and the level of interdependencies are comparatively low in the scheme of possible projects. However, this board game example is vastly different from the hardware technology projects that we study, where the numbers of components and interdependencies remain high across the different project types.

Figure A.2 shows how complexity can vary across projects in our sample. We highlight two examples from our sample: one project selected to exemplify what a low-complexity project looks like, and another strongly contrasting example of what high complexity looks like.

PROCESS FOR IDENTIFYING PROJECTS AND COMPLEXITY To identify these two examples, we systematically reviewed product images and descriptions on the campaign webpages. We isolate that sample to smaller and smaller subsamples based on the following criteria. First, the projects needed to provide enough images of the product to allow for a visual examination of the product. Second, the projects needed to have most of their features observable from the images. Projects which sophis-

ticated internal components were only considered if they provided images of the internal workings of the product. We also excluded projects where most of their capabilities were enabled by embedded software or firmware, which would not be observable to us. The ideal projects were the ones where most of their capabilities were exposed on the exterior of the product. Third, we focused on products where most of their value or innovation came from improvements in the fabrication of a custom circuit board to avoid examples that were not transparent to laypersons. For the projects that meet these criteria, we take discretionary liberties to focus on products that readers of the manuscript could appreciate the value and novelty of. We selected the two projects among that set that had among the least and most complexity based on the process described next. The low-complexity project is the Griffin Pocket Tool XL, a metal multi-purpose tool. The high-complexity project is the Obsidian 3D Printer.

We take a deep dive to document the complexity of these projects in as much detail as possible, describing the exact process that was used and subsequently independently verified by an expert hardware technology entrepreneur. First, as detailed as we can possibly observe from public records, we list out all the possible components or modules that make up the product. One can think of the set of components that generally reflect the feature set of the product. Each of these components is associated with several organizational tasks, which can happen within the focal entrepreneurial firm or by third-party suppliers to the entrepreneur.

We then list out the interdependencies that we know must exist relative to these product components and modules. These are just the ones that we can infer from a best-case scenario of how the product would be manufactured. To be clear, our effort in listing out these interdependencies is subject to the same bounded rationality limitations as the entrepreneur. To be even clearer, we are likely underestimating the total number of interdependencies by a substantial degree relative to the actual entrepreneur who obviously knows more about the product than we do. But the point we want to make is simple: the total number of interdependencies starts getting out of hand very rapidly. We are certain that our examination is incomplete and that we have left many “unforeseen interdependencies”

off our list.

IMPLICATIONS OF ILLUSTRATIVE PROJECTS From Figure A.2, there are two takeaways. First, the number of components or tasks, even on one of the simplest projects, is legitimately quite high already. Second, and more importantly, this results in an enormous number of interdependencies that need to be addressed, even for the simplest project. This suggests that we should infer that the general level of interdependencies, and even the lowest levels of interdependencies, is indeed quite high, and high enough to allow for interdependencies to increase rapidly as features are added.

A.1.3 ENTREPRENEURIAL PREDICTION PROCESS

We provide contextual background and supporting evidence for (1) the challenges inherent in making predictions and (2) the practical steps the entrepreneurs take to make these predictions.

WHAT PREDICTIONS ENTAIL

Through our qualitative interviews, we sought to understand the details of the prediction process undertaken by entrepreneurs, including the specific steps taken and items considered when predicting the project timeline.

Making predictions in any system with many interdependencies is very difficult. The entrepreneurs in our sample consistently recognize the complexity of the entrepreneurial endeavors they undertake. One noted, “The hardware game is hard. Even if you’ve been through... so many form factors and production issues, you think you would have perfected it. There is no perfecting manufacturing” (TabletCo CEO). Another highlighted the specific difficulties with launching hardware projects: “Software is very different if something goes wrong. You just push an update on the back end and it’s there in the morning. When you’re doing hardware and you miss a washer, that’s a huge freakin’ problem, and you’re for sure going to miss a washer sometime” (WidgetsCo CEO).

Our interviews suggest that the main source of prediction difficulty revolves around uncertainty relative to the interdependencies rather than the distinct components. The WidgetsCo CEO knew that the washer was a component of the project, but it was interdependencies with other parts of the project that could lead to the washer's omission that was the source of uncertainty. Entrepreneurs generally know all or most of the components required to complete the project, whereas a large number of the interactions between all the different components are unknown. In the language of the NK model (Kauffman, 1995), we interpret entrepreneurs making predictions in our context as having some reasonable sense for the value of N, but they also recognize that the bulk of the work in execution revolves around K, for which they have a much less accurate sense prior to actually executing the projects.

This is evident in our review of the entrepreneurs' own assessments of uncertainty, as they publicly report as "Risks and Challenges" required for every Kickstarter funding campaign. We review this content for a large set of projects and observe that the top risks that concern entrepreneurs relate to the interdependencies they might face in execution. They know the various components of the project (N) but don't yet know how these components fit together (K). One entrepreneur writes, "Every single component of this product is well known," but we "only have a 3D-printed prototype right now, and...don't have any experience with injection moulding at this point." Others note interdependencies inherent in "Coordination with multiple manufacturers, with various lead times"; "testing the pre-mass production sample device"; and obtaining "CE, FCC, IC compliance...certificates."

Part of this knowledge around components and uncertainty around interdependencies is true by construction in our setting. Kickstarter requires all entrepreneurs to have a working prototype. As a result, the entrepreneurs are already aware of the various components needed to produce the product. The complexity lies in the interdependencies that may not have surfaced when making the prototype: "The prototypes are all handmade. They're more of a unique product that has more time put into it. But when you're doing production, you're not spending that much time on every single unit. You're

doing large volume. That’s where we end up having problems” (GPSCo CEO).

HOW PREDICTIONS ARE MADE

Based on our qualitative interviews, we find that entrepreneurs generally make predictions by breaking down the project into specific tasks, making predictions relative to each specific task, and then aggregating the task-level predictions to the project level. CircuitsCo CEO described how, “I just took basically all the things I knew would take time—like waiting for Kickstarter to wire the money, manufacturing, shipping—--and I added all of them together.” 3DPrintCo CEO detailed a similar process for making the timeline prediction: “We’ve got a lot of experience in manufacturing. So it was a matter of understanding what our supply chain would look like and understanding how our product design could deviate from the design that we launched with so that whatever unforeseen challenges we may have faced we could have accomplished or overcome those challenges with deviations in our plan. So we tried to create basically a risk mitigation program for the possible design variations that we would have expected.”

Our interviews also support the notion that entrepreneurs are making these predictions for execution that follows a local search (Sommer & Loch, 2004). AccessoryCo CEO emphasized how his execution involved sequentially changing “one component after another” and “one iteration” at a time. For example, he first changed “the quality of the plastic, then the color of the plastic, ... then the adhesive tape.” This local search execution process is complementary to the process of aggregating the individual, step-by-step task-level predictions to the project level.

Once entrepreneurs determine their best project-level prediction, they seem to systematically “pad” their timelines with extra time as a precaution. CircuitsCo CEO said he usually “added like a month of buffer or something” and TabletCo CEO said he urges other entrepreneurs, “Don’t be too aggressive...definitely build in like two or three months extra.” This sentiment is echoed by WidgetsCo CEO who said, “Obviously, things go wrong. So another thing I do is I take my timeline and I add 30% to

it. I don't care what it is because you're going to fuck it up so yeah. So that's my rule for financial modeling and project modeling. Always add 30% because something's gonna go wrong.”

As detailed in Online Appendix Section A.1.10, we find no evidence that the prediction process for delivery timelines involves social comparison or competitive benchmarking with or against other entrepreneurs.

A.1.4 DATA COLLECTION

FEATURES RANK

To collect the *Features Rank* variable, we hired five independent reviewers to rank each entrepreneur's projects by the number of features.¹ We made two separate efforts to make sure this data generation process would be fruitful. First, to ensure a potential reviewer could handle what was asked of them, we carefully selected among reviewers to make sure they met a general qualification level. We sought out individuals with educational and professional experience in fields relevant to understanding and evaluating hardware products. Second, given that the projects span a relatively diverse set of product subtypes—although limited to only hardware products—it was important that we bring in a diverse set of reviewers such that, for any given project, the majority of reviewers would be qualified to make an assessment. We sought out reviewers of different genders and generations to obtain a balanced perspective. Table B.6 summarizes the backgrounds of the reviewers.

Each reviewer reviewed the photographs and product descriptions of each project by the same entrepreneur. We provided the reviewers with an Excel document containing sets of projects grouped by entrepreneur. The order of the projects within each entrepreneur group was randomized. The reviewers opened the links for the Kickstarter project page for each of the projects in the entrepreneur group. They then assigned a rank to each of the projects by the same entrepreneur. Even when projects seemed

¹The reviewers also ranked the projects according to technical sophistication, which is closely correlated to number of features.

Table A.3: Project Reviewer Backgrounds. Education and experience backgrounds of each of the five individuals hired to rank each entrepreneur’s set of projects according to number of features.

Reviewer	Education	Experience
Reviewer 1	MBA graduate with additional technical masters	12 years experience in computer programming
Reviewer 2	College degree in engineering	Career working at robotics company
Reviewer 3	College degree in business or engineering	2+ years work experience in consulting, banking, or engineering
Reviewer 4	Mechanical engineering and computer science	Freelance web design
Reviewer 5	Senior in high school	Experience in VC diligence and health technology

almost identical or very similar, we required the reviewers to force a ranking between all projects. Each reviewer repeated this process for all the groups of projects by the same entrepreneur.

We aggregate the rankings across the reviewers. In the case of disagreement between reviewers, we take the average rank between the two reviewers and then re-rank the projects based on the averaged scores. We allow for ties (which only occur in the case of conflicting rankings). Given that we force a ranking even when projects are nearly identical or very similar, this means that rank assignment is close to random in those circumstances. As a result, we expect some disagreement between reviewer rankings. This design gives more weight to the rankings that are more clear-cut and less weight to the rankings that are more ambiguous. If two projects tie for the most features, the binary indicator turns on for both projects, which washes out with the inclusion of entrepreneur fixed effects.²

DELAY DURATION AND ACTUAL TIME

To collect the actual shipment date for all the projects in our sample, we hired two contractors on UpWork.³ We provided them with an Excel file containing a link to each project and a row for each

²For example, if an entrepreneur has two total projects and the two reviewers disagree on which of the two projects has more features, then both projects would have an average rank of 1.5 (being rated as 1 by one reviewer and 2 by the other). After re-ranking based on these average scores, both of the entrepreneur’s projects would have a 1 for *Features Most*. Because we include entrepreneur fixed effects in the regression, this entrepreneur would show no change across projects. As a result, in our model the variation in number of features across an entrepreneur’s subsequent projects is driven by the projects where the ranking between projects is clearer and more objective and is not driven by the projects where the ranking is ambiguous.

³We use the date shipped as opposed to the date the product arrived on the customer’s doorstep. Shipment date is more standardized and consistent across all projects regardless of customer location. There are also gen-

of the rewards associated with that project. The contractor followed the link to the project page and read through the updates and comments to record (at the individual reward level) the date on which the reward first began to ship, the date on which all the rewards were shipped, whether it was apparent that the reward had still not shipped, or whether there was no mention of shipping. Across the 722 projects in our sample, there are 7,019 different rewards with an average of 9.7 rewards per project. We exclude the \$1 cash donation reward option included on most projects as well as rewards that were not backed. We were able to identify the date the reward started shipping for 71% of rewards and the date the reward finished shipping for 49% of rewards.

At the reward level, we take the later of the date the reward first shipped and the date the reward finished shipping. We take the later date because the entrepreneur's prediction is for the date all rewards will be delivered, not when the entrepreneur would deliver the first reward. This allows us to measure the actual shipment date for 80% of all rewards in our sample (which includes all rewards that list either the date the reward first shipped or the date the reward finished shipping). The actual delivery date has a day, month, and year whereas the predicted delivery date is just a month and a year. To be conservative, the delay for each reward is computed as the difference between the actual shipment date and the first day of the month after the month of the predicted delivery date. We then aggregate this reward-level data to the project level to compute the *Delay Duration* by taking the mean delay of all the rewards in each project. This same method of aggregating from the reward level to the project level is used for *Predicted Time*. The results are almost identical if we use the median. Using this approach we are able to measure the *Delay Indicator* for 89% of projects in the sample. This includes all projects with shipment information for at least one reward in the project.

erally only one or two shipment dates per project as opposed to many more unique dates when the backers receive the reward. And, as a practical matter, the data on when customers receive rewards is much sparser and more inconsistent than the data on when the rewards ship.

EXTERNAL FUNDING

To define our *External Funding* control variable, we collect data on venture financing from Crunchbase. We hired an RA to search both the entrepreneur name and the company name (if available) for each project to look for any matches on Crunchbase. For those entrepreneurs with a Crunchbase page, we collected data on each fundraising round (excluding those classified as crowdfunding, which would reflect the Kickstarter projects). This binary indicator turns on if the date the project launched comes after the date the entrepreneur of that project raised external capital from another source. The results do not change in any substantive way if we use the cumulative number of external funding rounds instead of this binary indicator.

A.1.5 VARIABLE DISTRIBUTIONS

QUANTILE SUMMARY STATISTICS

For transparency and clarity, we supplement the summary statistics provided in Table 1.1. Table A.4 includes the 0%, 25%, 50%, 75%, and 100% quantiles for each independent variable used in analysis.

DISTRIBUTION VISUALIZATION

In addition, we include a visualization of the distribution of the variables. Figure A.3 contains plots of the distribution of each independent variables used in analysis. Density plots are used for continuous measures, and histograms are used for measures that fall into a small number of finite values or categories.

Table A.4: Variable Quantiles. 722 project-level observations. *Funding Threshold*, *Funding Exceeded*, and *Funding Backers* are all in thousands.

Variables	Quantiles				
	0%	25%	50%	75%	100%
Project Experience	0	0	1	1	5
Failed Campaign Experience	0	0	0	0	4
Prior Project Funding Deviation	-1	0	0.05	2.94	85.65
Prior Project Delay	-0.53	0	0	0.42	10.78
Execution Overlap	0	0	0	0	1
External Financing	0	0	0	0	1
New Category	0	0	0	0	1
Elapsed Time	0	0	112.50	538.75	2,458
Prior Updates	0	3	6	9	40
Funding Period	2	30	30	37	60
Funding Reward Tiers	1	6	9	12	34
Funding Reward Size	4	50	99	201.63	5,995
Funding Threshold	0.02	3.85	10.15	30	261.96
Funding Exceeded	0	3.51	17.59	74.92	3,351.36
Funding Backers	0.001	0.12	0.30	0.92	28.14

A.1.6 UNFORESEEN INTERDEPENDENCIES ALTERNATIVE

We find that all results hold if, instead of using *Unforeseen Interdependencies*, we define a new variable, *Alternative Unforeseen Interdependencies*, which is equal to the count of updates that contain words from *both* sets of relevant words. As shown in Table A.5, we find that this alternative definition of *Alternative Unforeseen Interdependencies* ($p \sim 0.002$) is still positively related to *Project Experience*. On each subsequent project, entrepreneurs disclose encountering 0.114 additional unforeseen interdependencies that contain references to both being unforeseen and dealing with interdependencies. Using this alternative definition, \ln *Alternative Unforeseen Interdependencies* ($p = 0.001$) also positively associates with *Project Experience*. Each subsequent project increases unforeseen interdependencies by 6.4%. The effect size is smaller using this alternative definition given that requiring updates to contain words from both sets is a stricter criterion than including updates with words from either set and, as a result, fewer total updates meet this criterion.

Table A.5: Unforeseen Interdependencies Alternative. Ordinary least squares (OLS) estimation. Robust standard errors are clustered at the entrepreneur level. p -values are shown in parentheses.

	Alternative Unforeseen Interdependencies	Ln Alternative Unforeseen Interdependencies
Project Experience	0.114 (0.002)	0.064 (0.001)
Failed Campaign Experience	0.052 (0.369)	0.030 (0.422)
Prior Campaign Funding Deviation	-0.008 (0.068)	-0.003 (0.051)
Prior Project Delay	-0.045 (0.041)	-0.029 (0.032)
Execution Overlap	0.022 (0.788)	0.012 (0.756)
External Financing	-0.120 (0.620)	-0.085 (0.549)
New Category	0.124 (0.112)	0.064 (0.069)
Elapsed Time	0.000 (0.893)	0.000 (0.366)
Funding Period	0.001 (0.872)	0.002 (0.236)
Funding Reward Tiers	0.003 (0.559)	0.001 (0.689)
Funding Reward Size	-0.000 (0.991)	-0.000 (0.914)
Ln Funding Threshold	0.047 (0.161)	0.018 (0.289)
Ln Funding Exceeded	0.090 (0.009)	0.050 (0.011)
Ln Funding Backers	-0.022 (0.753)	-0.030 (0.292)
Baseline Updates	-0.004 (0.537)	
Ln Baseline Updates		-0.005 (0.796)
Entrepreneur FE	Yes	Yes
Product Subtype FE	Yes	Yes
Year FE	Yes	Yes
Month FE	Yes	Yes
R ²	0.502	0.560
Entrepreneurs	314	314
Observations	722	722

A.1.7 INCREASING FEATURES, UNFORESEEN INTERDEPENDENCIES, AND DELAY

We first test the assumption that, as more features are added to a project, the unforeseen interdependencies will increase, showing a positive relationship between an increase in the number of features and an increase in unforeseen interdependencies.

We define *Features Increase* as a binary indicator of whether the entrepreneurs add more features for their next projects, which is equal to 1 if the current project has more features than the prior project and 0 otherwise. We use *Features Increase* as the main independent variable for this analysis and look at the relationship between it and three different measures of whether the unforeseen interdependencies will increase as dependent variables. *Unforeseen Interdependencies Increase Binary* is a binary indicator equal to 1 if the current project has more updates related to unforeseen interdependencies than the prior project. *Unforeseen Interdependencies Increase Count* is equal to the number of additional updates related to unforeseen interdependencies than the prior project. *Unforeseen Interdependencies Increase IHS* takes the inverse hyperbolic spline of *Unforeseen Interdependencies Increase Count* in order to account for any potential non-linearities in the relationship and to reduce the impact of any outliers. We find a significant and positive relationship between *Features Increase* and each of our three measures of increasing unforeseen interdependencies. If features are added to the project, the project is 27% more likely to have more unforeseen interdependencies than the prior project or on an absolute basis will encounter 1.721 more unforeseen interdependencies (an increase of 0.735 when taking the inverse hyperbolic spline).

We then test the assumption that, as more features are added to a project, the delay will increase, showing a positive relationship between an increase in the number of features and an increase in delay. We use the same independent variable, *Features Increase*, as defined above. We look at the relationship between this independent variable and two measures of delay, *Delay Indicator* and *Delay Duration* as defined in the main paper. We use all the same controls and fixed effects as specified in the main paper.

We find that, if features are added to a project, that project is 9.4% more likely to be delayed and, on average, will be delayed by 18.298 additional days.

A.1.8 NON-LINEAR EFFECT OF EXPERIENCE

We find empirical evidence for a faster-than-linear increase in unforeseen interdependencies and delay duration. This empirical evidence suggests that we more than meet the minimum set of assumptions required by our theory, e.g., the project complexity curve increases faster than a concave learning curve. Figure 2.3 of the main manuscript is the visual representation of the estimates from the regression models we describe next.

We considered several econometric models—including those that specify a specific functional form of the relationship—but we decided that it would be best to remain agnostic to functional form. Instead, we construct a series of indicator variables representing different levels of *Project Experience* that allow us to flexibly and non-parametrically identify the relationship. This type of specification allows the functional form to “reveal” itself to us without us having to pre-specify its exact shape.

In place of *Project Experience*, we include a set of indicator variables. *Project Experience: Second* takes a value of 1 if the focal project is the entrepreneur’s second project, and 0 otherwise. *Project Experience: Third* and *Project Experience: Fourth or More* follow similarly. We group together experience for entrepreneurs on their fourth or later project because the number of observations on these higher levels of experience is quite sparse and thus noisy; the observed pattern of statistically significant results is robust to the exclusion of projects that are the entrepreneur’s fifth or later project. Since we omit the indicator for an entrepreneur’s first project (zero experience), all coefficients on these three indicators should be interpreted as relative to the scenario of the entrepreneur’s first project. Table 3.3 and Table 2.6 are replicated in Table A.6 and Table A.7, respectively, swapping out *Project Experience* for this set of indicator variables.

Table A.6: Non-Linear Analysis of Complexity and Unforeseen Interdependencies. Ordinary least squares (OLS) estimation. Robust standard errors are clustered at the entrepreneur level. *p*-values are shown in parentheses.

	Features Most	Features Rank	Features Percentile	Unforeseen Interdependencies	Ln Unforeseen Interdependencies
Project Experience: Second	0.266 (0.001)	0.417 (0.000)	0.302 (0.000)	1.457 (0.000)	0.270 (0.001)
Project Experience: Third	0.188 (0.133)	0.541 (0.011)	0.234 (0.048)	2.721 (0.000)	0.420 (0.002)
Project Experience: Fourth or More	0.508 (0.024)	1.468 (0.000)	0.457 (0.026)	4.116 (0.000)	0.738 (0.003)
Failed Campaign Experience	0.091 (0.432)	0.184 (0.214)	0.103 (0.291)	0.632 (0.086)	0.184 (0.056)
Prior Campaign Funding Deviation	-0.001 (0.831)	0.000 (0.984)	-0.002 (0.653)	0.003 (0.888)	-0.001 (0.740)
Prior Project Delay	-0.004 (0.917)	-0.038 (0.475)	-0.015 (0.665)	-0.710 (0.000)	-0.133 (0.000)
Execution Overlap	-0.440 (0.005)	-0.488 (0.012)	-0.412 (0.006)	-0.295 (0.651)	0.063 (0.591)
External Financing	0.332 (0.010)	0.244 (0.113)	0.321 (0.011)	1.087 (0.336)	0.135 (0.559)
New Category	-0.013 (0.937)	0.179 (0.450)	0.075 (0.663)	1.686 (0.074)	0.272 (0.211)
Elapsed Time	-0.000 (0.435)	0.000 (0.900)	-0.000 (0.776)	0.004 (0.009)	0.001 (0.000)
Funding Period	0.005 (0.228)	0.008 (0.141)	0.005 (0.158)	0.028 (0.189)	0.006 (0.196)
Funding Reward Tiers	0.007 (0.345)	0.008 (0.505)	0.008 (0.254)	-0.024 (0.502)	-0.004 (0.634)
Funding Reward Size	0.000 (0.002)	0.000 (0.000)	0.000 (0.001)	-0.000 (0.796)	0.000 (0.814)
Ln Funding Threshold	0.035 (0.429)	0.043 (0.424)	0.028 (0.503)	0.745 (0.004)	0.156 (0.000)
Ln Funding Exceeded				0.679 (0.008)	0.080 (0.093)
Ln Funding Backers				-0.032 (0.929)	0.069 (0.299)
Baseline Updates				0.175 (0.001)	
Ln Baseline Updates					0.283 (0.000)
Entrepreneur FE	Yes	Yes	Yes	Yes	Yes
Product Subtype FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes
R ²	0.301	0.442	0.299	0.745	0.775
Entrepreneurs	314	314	314	314	314
Observations	722	722	722	722	722

Table A.7: Non-Linear Analysis of Delivery and Delay. Ordinary least squares (OLS) estimation. Robust standard errors are clustered at the entrepreneur level. p -values are shown in parentheses.

	Delay Indicator	Delay Duration	Delay Percent	Predicted Time	Actual Time
Project Experience: Second	0.115 (0.038)	44.858 (0.003)	0.717 (0.000)	6.082 (0.151)	49.112 (0.004)
Project Experience: Third	0.262 (0.013)	75.369 (0.006)	1.109 (0.001)	18.816 (0.010)	89.149 (0.004)
Project Experience: Fourth or More	0.362 (0.010)	145.106 (0.002)	1.581 (0.001)	30.063 (0.028)	169.153 (0.001)
Failed Campaign Experience	0.134 (0.132)	-5.134 (0.766)	-0.053 (0.810)	14.622 (0.047)	7.842 (0.739)
Prior Campaign Funding Deviation	-0.003 (0.074)	0.034 (0.939)	-0.004 (0.620)	0.127 (0.450)	0.361 (0.464)
Prior Project Delay	-0.041 (0.050)	-21.968 (0.023)	-0.129 (0.640)	4.629 (0.021)	-19.659 (0.037)
Execution Overlap	-0.050 (0.474)	8.015 (0.839)	1.124 (0.248)	-3.411 (0.687)	6.427 (0.883)
External Financing	0.140 (0.198)	91.972 (0.254)	0.317 (0.548)	22.198 (0.358)	121.857 (0.217)
New Category	0.086 (0.596)	41.691 (0.109)	0.477 (0.091)	5.572 (0.576)	48.525 (0.149)
Elapsed Time	-0.000 (0.448)	-0.096 (0.100)	-0.001 (0.285)	-0.022 (0.271)	-0.112 (0.078)
Funding Period	0.009 (0.000)	1.272 (0.072)	0.003 (0.899)	0.190 (0.448)	1.739 (0.023)
Funding Reward Tiers	0.003 (0.479)	0.765 (0.489)	0.000 (0.989)	-0.165 (0.692)	0.968 (0.459)
Funding Reward Size	0.000 (0.792)	-0.011 (0.707)	0.000 (0.680)	0.005 (0.525)	-0.009 (0.814)
Ln Funding Threshold	0.006 (0.842)	22.322 (0.107)	-0.081 (0.582)	19.326 (0.000)	40.085 (0.011)
Ln Funding Exceeded	0.035 (0.194)	29.877 (0.003)	0.205 (0.081)		29.323 (0.010)
Ln Funding Backers	0.007 (0.852)	-35.697 (0.006)	-0.183 (0.265)		-33.133 (0.031)
Entrepreneur FE	Yes	Yes	Yes	Yes	Yes
Product Subtype FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes
R ²	0.663	0.719	0.657	0.753	0.738
Entrepreneurs	306	303	303	314	303
Observations	686	644	644	722	644

We outline the effects for a subset of the key variables here, with the full results shown in the regression tables.

FEATURES

All relative to the entrepreneur's first project, the average entrepreneur's second project has 0.417 higher feature rank, the third project has an average of 0.541 higher feature rank, and the fourth or later project has 1.468 higher feature rank. To give an idea of the comparison across projects, we can look at the difference between these coefficient values, with the second project having a feature rank 0.417 higher than the first project, the third project having a feature rank 0.124 higher than the second project, and the fourth or later project having a feature rank 0.927 higher than the third project. We also look at the same interpretation for unforeseen interdependencies.

UNFORESEEN INTERDEPENDENCIES

All relative to the entrepreneur's first project, the average entrepreneur's second project encounters 1.457 more unforeseen interdependencies, the third project encounters an average of 2.721 more unforeseen interdependencies, and the fourth or later project encounters 4.116 more unforeseen interdependencies. To give an idea of the comparison across projects, we can look at the difference between these coefficient values, with the second project encountering 1.457 more unforeseen interdependencies than the first project, the third project encountering 1.264 more unforeseen interdependencies than the second project, and the fourth or later project encountering 1.395 more unforeseen interdependencies than the third project.

DELAY DURATION

All relative to the entrepreneur's first project, the average entrepreneur's second project is delayed by an additional 44.858 days, the third project is delayed by an additional 75.369 days, and the fourth or later project is delayed by an additional 145.106 days. To give an idea of the comparison across projects, we can look at the difference between these coefficient values, with the second project delayed by 44.858 more days than the first project, the third project delayed by 30.511 more days than the second project, and the fourth or later project delayed by 69.737 more days than the third project.

Furthermore, to provide readers with more intuition on the holistic pattern implied by the above regression estimates, we generate a visualization of the estimates for the effect of various levels of project experience on *Actual Time* and *Predicted Time*. In some sense, these variables intuitively map to actual project complexity and predicted project complexity curves, respectively. Figure 2.3 of the main manuscript plots the coefficients, with the project number on the horizontal axis and the *Actual Time* and *Predicted Time* (both in days) on the vertical axes.⁴ This figure shows that the actual delivery time increases much more sharply relative to the predicted delivery time, with the gap between actual delivery time and predicted delivery time increasing as entrepreneurs gain experience.

A.1.9 PRIOR CAMPAIGN FUNDING AND BEHAVIOR

THEORETICAL BACKGROUND

We explore in depth whether prior campaign funding outcomes impact an entrepreneur's behavior on subsequent projects. We do not intend to make any groundbreaking theoretical claims on this point: our primary goal is to make sure we properly account for and apply classic behavioral theory

⁴To set the level of the omitted coefficient of the entrepreneur's first project—and thus the level of all the estimates as they are relative to that baseline—we calculate the mean value of each variable used in the regression and multiply that average value by the corresponding coefficient. We sum those values and then add the mean entrepreneur fixed effect.

on performance feedback and outcome–aspiration gaps (Cyert & March, 1963, Greve, 1998). We apply this theory to product introductions (Joseph & Gaba, 2015). In short, we consider the theoretical argument that when an entrepreneur suffers from an outcome–aspiration gap in their prior experience—specifically, they suffer from a failed funding campaign on the previous project—the entrepreneur would have higher risk tolerance and engage in problemistic search on the next project (Greve, 2003). Assuming that this feedback is sufficiently unambiguous to trigger the entrepreneur to respond (Joseph & Gaba, 2015), the entrepreneur would take on a project that has a greater likelihood of unforeseen complexity that could delay the project. As shown in Figure A.4, this relationship can be visualized as a “V”, where the outcome–aspiration gap is on the x-axis and the accuracy of strategic foresight is on the y-axis. This “V” shape would manifest if these were a linear relationship; however, if the relationship is non-linear, we would see more of a “U” shape, which is also included in the visualization.

Recent research adds additional nuance to the theory that would strengthen the argument. Keum & Eggers (2018) argue that managers would set more aggressive aspirations, like on project complexity and timeline, when facing increased pressure to acquire resources, like if they feared missing funding targets on the next project because they missed them on the prior project. Eggers & Kaul (2018) argue that firms over-invest in radical invention when performance is moderately below aspiration, whereas in our setting a radical invention would be a radical departure from their previous project (which would add a lot of new complexity).

EMPIRICAL CONTEXT

Turning to our specific empirical context, it is quite rare for these serial-project entrepreneurs to fail in their fundraising efforts. Of the 314 entrepreneurs who completed successful projects in our sample, only 33 entrepreneurs had previously run a failed funding campaign (10.5% of entrepreneurs), with 42 failed funding campaigns total out of 782 total funding campaigns (5.4% of funding campaigns).

At a project level, our level of analysis, only 36 (5.0%) out of 722 projects in our sample (i.e., the successful projects) had a failed funding campaign of the same product subtype immediately prior to the focal project. In comparison, Kickstarter reports that 61.6% of all projects fail in their funding campaigns, and specifically 79.0% of technology projects (where our sample originates) fail in their funding campaigns.⁵

Clearly, there is a compositional difference between our sample of projects and entrepreneurs and the universe present on Kickstarter. First, we study serial-project entrepreneurs, a more professional set of entrepreneurs who tend to treat their projects as full-time jobs; in many cases, there are entrepreneurial firms behind the effort. Second, we focus on technical hardware product categories, of which there might be more consumer interest and that may have a higher barrier to entry, i.e., it takes a significant amount of effort to even create the prototype that gets presented on the fundraising page.

Thus, our sample of entrepreneurs may not be the best sample on which to study the implications of prior funding failure as a general phenomenon. Nevertheless, we proceed with an empirical exploration that accounts for this past project funding failure (and success).

VARIABLES

As the main independent variable capturing the outcome–aspiration gap, we use *Prior Campaign Funding Deviation*, which is equal to the percentage by which the prior funding campaign exceeded (or missed) its funding goal. Values of *Prior Campaign Funding Deviation* less than 0 occur when the entrepreneur failed to meet her prior campaign’s funding threshold, and greater than or equal to 0 occur when the entrepreneur succeeded. We interact *Prior Campaign Funding Deviation* with an indicator variable *Prior Campaign Funding Success* that takes a value of 1 if *Prior Campaign Funding Deviation* is greater than or equal to 0, and 0 otherwise. This interaction term allows us to estimate separate slopes for the two halves of the “V” shape outlined in the theoretical background.

⁵<https://www.kickstarter.com/help/stats>, accessed December 2020.

STATISTICAL MODEL

We need to also consider that the effect of performance feedback may be heterogeneous and non-linear, e.g., greater degrees of success or failure have a greater effect size per unit of deviation than success or failure that is close to the aspiration level. In other words, a U-shaped relationship rather than a V-shaped relationship. Thus, we also test a model that enters in a quadratic term for *Prior Campaign Funding Deviation*, and we interact both the base term and the quadratic term with *Prior Campaign Funding Success* to allow estimates of different “curves” above and below the aspiration level. In theory, this would allow us to estimate the two halves of a theoretical “U.”

DESCRIPTIVE VISUALIZATION

Before we turn to the regression analysis, we generate descriptive plots of the relationship between *Prior Campaign Funding Deviation* and our two main measures of performance: (i) *Delay Percent* (defined as *Delay Duration* divided by *Predicted Time*) and (ii) *Unforeseen Interdependencies*. When generating the plots, we exclude outliers for clarity in visualization, but all trends and interpretations hold when including outliers. Examining these plots in Figure A.5 and Figure A.6, we see similar trends in both plots, with a positive slope where *Prior Campaign Funding Deviation* is less than zero (though with a very wide confidence interval given the very limited number of observations) and a flat or very slightly increasing trend where *Prior Campaign Funding Deviation* is greater than zero (though again with an increasingly wide confidence interval moving away from the bulk of the data).

RESULTS

Table A.8 shows this relationship when including controls and fixed effects. Given the focus on the lagged funding deviation, we exclude first projects where there is no defined lagged funding deviation. As a result, we also exclude entrepreneur fixed effects which are not appropriate for entrepreneurs

with only a single project after their first project, leaving no variation within the entrepreneur's set of projects. As suggested by the visual evidence, we do not observe any significant relationship between *Prior Campaign Funding Deviation* and *Delay Percent* or *Unforeseen Interdependencies*. The directionality of the point estimates matches the visual evidence, with a positive coefficient on *Prior Campaign Funding Deviation*, suggesting a positive slope when *Prior Campaign Funding Success* is equal to zero (the area to the left of zero on the plots) and then a flat slope when *Prior Campaign Funding Success* is equal to one and the coefficients are summed to give a slope around zero. We also find no significant relationships when including a quadratic term.

Given the small sample size of projects that missed their funding goal, we are limited in our ability to interpret the trend where *Prior Campaign Funding Success* is equal to zero. We do have ample data where *Prior Campaign Funding Success* is greater than zero, but again we find no significant relationship. One possible explanation is that the impact of deviation from the funding target on the prior project is overshadowed by the impact of deviation from the funding target on the current project (which is included as a control). Another possible explanation is that *Delay Percent* and *Unforeseen Interdependencies* are both measures of execution and are therefore one step removed from the impacts of fundraising outcomes.

A.1.10 SOCIAL COMPARISON

We seek to address potential ambiguity around what could be driving prediction failure in this setting. If the entrepreneur believes that setting a delivery timeline comparable to her peers is important to fundraising, the prediction failures could be interpreted as a matter of entrepreneurs socially informed about their competitive context. To test, and ultimately rule out, this alternative explanation, we define the entrepreneur's peer group and then empirically test the impact of the peer group on the entrepreneur's predicted time as well as the impact of deviation from the peer group on ability to fundraise.

Table A.8: Funding Deviation and Performance. Ordinary least squares (OLS) estimation. Robust standard errors are clustered at the entrepreneur level. p -values are shown in parentheses.

	Delay Percent		Unforeseen Interdependencies	
Prior Campaign Funding Deviation	0.519 (0.569)	-7.227 (0.356)	1.249 (0.468)	5.071 (0.498)
Prior Project Funding Success	-0.110 (0.898)	1.665 (0.247)	0.991 (0.420)	0.108 (0.954)
Prior Campaign Funding Deviation \times Prior Project Funding Success	-0.524 (0.565)	7.238 (0.356)	-1.260 (0.465)	-5.010 (0.503)
Prior Campaign Funding Deviation ²		-6.814 (0.351)		3.835 (0.572)
Prior Campaign Funding Deviation ² \times Prior Project Funding Success		6.814 (0.351)		-3.836 (0.572)
Project Experience	-0.045 (0.806)	-0.037 (0.841)	-0.106 (0.769)	-0.136 (0.713)
Failed Campaign Experience	0.081 (0.736)	0.135 (0.580)	0.647 (0.230)	0.651 (0.228)
Prior Project Delay	0.662 (0.023)	0.663 (0.022)	0.193 (0.238)	0.203 (0.222)
Execution Overlap	1.431 (0.210)	1.427 (0.213)	0.734 (0.168)	0.765 (0.155)
External Financing	-0.238 (0.550)	-0.264 (0.500)	0.917 (0.167)	0.799 (0.228)
New Category	-0.765 (0.196)	-0.640 (0.263)	2.876 (0.100)	3.213 (0.072)
Elapsed Time	-0.000 (0.277)	-0.000 (0.285)	0.000 (0.987)	0.000 (0.955)
Funding Period	0.012 (0.343)	0.012 (0.328)	0.003 (0.875)	0.005 (0.790)
Funding Reward Tiers	0.008 (0.565)	0.009 (0.515)	0.077 (0.053)	0.083 (0.042)
Funding Reward Size	0.000 (0.303)	0.000 (0.309)	0.000 (0.618)	0.000 (0.630)
Ln Funding Threshold	-0.069 (0.440)	-0.065 (0.473)	0.762 (0.000)	0.761 (0.000)
Product Subtype FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
R ²	0.205	0.206	0.262	0.271
Sample	2+ Proj	2+ Proj	2+ Proj	2+ Proj
Observations	342	342	402	402

DEFINING ENTREPRENEUR PEER GROUP

We consider two dimensions when defining an entrepreneur's peer group that may impact their behavior through social comparison. First, we define a set of comparable projects the entrepreneur could reasonably view as competition. At a reductive level, all Kickstarter projects that are soliciting the same dollars are in competition. The most competitive set of projects seem to be those within same product subtype, e.g., 3D printing, camera equipment, wearables. The key assumption we make here is that the entrepreneur perceives that she is competing with those projects, based on our assumption of the entrepreneur's assumption that her customers navigate and search through Kickstarter by product subtype. That said, based on our interviews, entrepreneurs do not seem to view competition for crowdfunding as a zero-sum game, given that the vast majority of customer spending is not on Kickstarter projects and that many backers find their way directly to a project without navigating through the Kickstarter platform, e.g., by a direct link from an organic social media campaign or direct-response online advertising.

Second, we define a time window during which the entrepreneur could reasonably have taken into account comparable projects prior to specifying her prediction for her own product specification and delivery date. As first order, we should only include projects prior to the launch of the focal campaign. We decided to include only peer projects that successfully completed their funding prior to the focal campaign; it seems unlikely the entrepreneur would benchmark herself against failed funding campaigns. In addition, intuitively it seems unlikely the entrepreneur would search deeply into the distant past to benchmark herself: more recent projects likely matter more since they reflect the current state of the market the entrepreneur would face. We set a threshold of one year, meaning that we only include peer funding campaigns launched within one year of the focal campaign. In summary, we specify the entrepreneur's peer group as the projects within the last year by other entrepreneurs in the sample that successfully met their funding threshold.

SETTING DELIVERY TIME

We consider whether the peer group timeline has an effect on the predicted timeline set by the entrepreneur. The dependent variable *Predicted Time* is the time in days between the end of the fundraising campaign and the predicted delivery date. Using the definition of the peer group previously explained, the main independent variable *Peer Group Predicted Time* is the average *Predicted Time* across all projects in the focal project's peer group.⁶

RESULTS As shown in Table A.9, we find no statistically significant relationship between *Peer Group Predicted Time* and *Predicted Time*. The coefficient and statistical significance of *Project Experience* remains consistent. These findings align with our understanding of how entrepreneurs in our study actually set their project timelines, based on our qualitative interviews. To summarize our understanding of this process, which is outlined in more detail in Online Appendix Section A.1.3, entrepreneurs seem to be giving their best estimate of how long they believe the project will actually take—summing the estimated time for each project component—and then adding some buffer time on top of their best guess at the predicted time. No entrepreneur in any of our interviews mentioned benchmarking their predicted time against the predicted time of other projects, or trying to game the predicted time in any other way.

INCENTIVE ALIGNMENT

We also explore whether the difference between an entrepreneur's predicted time from the average predicted timeline of other comparable projects impacts the amount of money the entrepreneur is able to raise. If estimating a shorter predicted time has pecuniary benefits in the fundraising process,

⁶For example, one entrepreneur in the sample launched a hardware project on October 11, 2018. His peer group contains all hardware projects launched prior to October 11, 2018 but after October 11, 2017 (one year prior). This defines a set of 11 projects with predicted times ranging from 10 days to 149 days. Taking the average predicted time across all 11 projects gives a *Peer Group Predicted Time* of 89 days.

Table A.9: Predicted Time and Peer Group Comparison. Ordinary least squares (OLS) estimation. Robust standard errors are clustered at the entrepreneur level. *p*-values are shown in parentheses.

	Predicted Time
Project Experience	7.671 (0.031)
Peer Group Predicted Time	0.042 (0.736)
Failed Campaign Experience	18.161 (0.030)
Prior Campaign Funding Deviation	0.169 (0.304)
Prior Project Delay	4.035 (0.028)
Execution Overlap	-4.249 (0.608)
External Financing	14.281 (0.589)
New Category	9.998 (0.310)
Elapsed Time	-0.017 (0.435)
Funding Period	0.229 (0.356)
Funding Reward Tiers	-0.089 (0.834)
Funding Reward Size	0.004 (0.583)
Ln Funding Threshold	19.965 (0.000)
Entrepreneur FE	Yes
Product Subtype FE	Yes
Year FE	Yes
Month FE	Yes
R ²	0.757
Entrepreneurs	314
Observations	712

entrepreneurs would be incentivized to benchmark against their peer projects. If this is not the case, that would imply that exceeding the funding threshold would not rely on benchmarking against a social comparison.

In this analysis, and as alluded to above, the main independent variable *Peer Group Deviation* is the difference between the *Predicted Time* of the focal project and the *Peer Group Predicted Time* of the focal project. We examine the relationship between *Peer Group Deviation* and two measures of exceeding the funding threshold. *Funding Exceeded* is equal to dollars raised in the focal project in excess of the *Funding Threshold*. *Funding Positive Deviation* is the percentage by which the funding threshold was exceeded, equal to *Funding Exceeded* divided by *Funding Threshold*. Because *Funding Threshold* is used in the derivation of both *Funding Exceeded* and *Funding Positive Deviation*, we include regressions including and excluding it as a control.

RESULTS As shown in Table A.10, we see a mixture of significant and insignificant relationships between *Peer Group Deviation* and our measures of exceeding the funding threshold. However, in the cases where the result is significant, the value is small and positive. This would suggest that entrepreneurs would be incentivized to give themselves slightly more time than the average among their peers, which is the exact opposite of the narrative that increasing delays are due to pressure to predict shorter delivery times relative to the peer group. Together, these regressions provide insignificant or contradictory evidence of the alternative explanation.

While entrepreneurs are incentivized to maximize funds raised, they also recognize and experience real pecuniary consequences for failing to meet their predicted time, as outlined in Online Appendix Section A.1.1.1. As a result, even if there were some benefits to predicting shorter delivery times, entrepreneurs that we interviewed seemed unwilling to make the tradeoff of estimating a shorter delivery time now with the expectation to need to delay later. For example, one notes that, “Whatever goodwill you built up beforehand, it’s like so discounted by the time you have to announce delays” (TabletCo

Table A.10: Funding and Peer Group Comparison. Ordinary least squares (OLS) estimation. Robust standard errors are clustered at the entrepreneur level. *p*-values are shown in parentheses.

	Ln Funding Exceeded		Funding Positive Deviation	
Peer Group Deviation	0.002 (0.131)	0.002 (0.085)	0.019 (0.013)	-0.002 (0.822)
Project Experience	0.045 (0.668)	0.049 (0.644)	1.792 (0.009)	1.394 (0.053)
Prior Campaign Funding Deviation	-0.155 (0.306)	-0.167 (0.262)	-2.693 (0.002)	-1.204 (0.100)
Failed Campaign Experience	-0.024 (0.000)	-0.024 (0.000)	-0.422 (0.000)	-0.449 (0.000)
Prior Project Delay	0.039 (0.498)	0.037 (0.525)	0.132 (0.622)	0.367 (0.189)
Execution Overlap	-0.363 (0.113)	-0.376 (0.094)	0.645 (0.638)	2.313 (0.190)
External Financing	0.835 (0.286)	0.838 (0.282)	2.277 (0.264)	1.914 (0.327)
New Category	-0.098 (0.739)	-0.110 (0.709)	0.783 (0.606)	2.279 (0.258)
Elapsed Time	0.001 (0.417)	0.001 (0.417)	0.001 (0.834)	0.003 (0.677)
Funding Period	0.016 (0.002)	0.016 (0.002)	0.094 (0.004)	0.067 (0.047)
Funding Reward Tiers	0.030 (0.036)	0.031 (0.030)	0.126 (0.103)	0.019 (0.818)
Funding Reward Size	0.000 (0.219)	0.000 (0.195)	0.002 (0.040)	0.000 (0.897)
Ln Funding Threshold	0.032 (0.671)		-4.043 (0.000)	
Entrepreneur FE	Yes	Yes	Yes	Yes
Product Subtype FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
R ²	0.850	0.850	0.813	0.783
Entrepreneurs	314	314	314	314
Observations	712	712	712	712

CEO).

A.1.11 LEARNING THAT DELAY IS “ACCEPTABLE”

To test whether entrepreneurs could learn that a delay is acceptable, we need to test whether entrepreneurs face consequences for delay. Li & Martin (2019) study this exact question in the Kickstarter context. Specifically, they look at the impact of failing to meet predictions on the entrepreneur’s reputation and subsequent ability to raise money. Importantly, they find that, all else equal, if an entrepreneur defaults on what they promised, the probability that they reach their funding goal on their subsequent project drops by 50%. The key mechanism for reputation formation is the project comments left by investors. As a result, they conclude that “entrepreneurs likely have incentives to deliver the product or service they promised as long as the backers have the ability to provide product/service feedback to the public.”

Given that public comments are a demonstrated mechanism for reputation formation and, as a result, present real, pecuniary consequences to the entrepreneur, we empirically measure the relationship between delay and public comments in our sample. Specifically, we find a positive relationship between *Delay Duration* and both *Total Comments* ($p \sim 0.002$) and *Negative Comments* ($p = 0.000$). To calculate comment sentiment, we use a standard R package (Rinker, 2019) to calculate the sentiment of each of the most recent 100 project comments (excluding comments by the entrepreneur). We then sum the number of negative comments for each project. As summarized in Table A.11, for each additional day of delay, there are 1.7 additional comments and the number of negative comments increases by 0.016. This suggests that entrepreneurs do experience consequences when they delay and are, therefore, incentivized to deliver on time. In addition, to allow for the possible moderating effect of *Project Experience* on the effect of *Delay Duration* on the *Negative Comments* generated by customers, we also include a model with an interaction term for *Delay Duration* and *Project Experience*. We find that the interaction term has a significant and positive coefficient. Overall, we find no

evidence, in terms of their public feedback, that customers are less concerned with a project delay because the entrepreneur is more experienced. In fact, the customers seem to penalize the entrepreneur with more negative comments on the entrepreneur's subsequent projects.

However, it could be the case that customers are more lenient towards delay for projects with more features. If this is the case, and the entrepreneur becomes aware of this fact as they gain experience, they may learn that delay is acceptable particularly for projects with more features. To test whether this is the case, we analyze the relationship between number of negative comments (*Negative Comments*) and the complexity of the product in terms of features: *Features Most*, *Features Rank*, *Features Percentile*. We include models both with standalone terms for these measures of features as well as a model for each features measure interacted with *Delay Duration*, with the dependent variable of *Negative Comments*. Note that *Negative Comments* are roughly proportional to *Delay Duration*, given that there is more time for customers to enter negative feedback the longer the delay (in a hazard model sense) and that longer delays would of course agitate customers more. We interact the measures of features with *Delay Duration* to explore whether there is any heterogeneity in the possible customer leniency mechanism, e.g., the interaction term would be significant and negative if, at higher levels of product complexity, the customers made fewer *Negative Comments*, particularly under longer *Delay Duration* situations. As reported in Table A.12, we find that the interaction term between *Delay Duration* and all three measures of features is insignificant, as well as negligibly small.

Our interviews with entrepreneurs confirm these findings. One entrepreneur notes, "There is definitely some pressure to deliver things on time...There definitely is pressure making sure things are right, making sure you don't have to take additional steps" (CircuitsCo CEO). And entrepreneurs do believe that there are consequences to delay: "[When a project gets delayed,] people really bash the product on review channels and pages and the comments section" (TabletCo CEO). Or, in the words of another, "[When a project gets delayed,] you get angry people. For us, image and branding is the most important thing, and you want to strive to deliver a really good product....and it does hurt your

Table A.11: Consequences of Delay. Ordinary least squares (OLS) estimation. Robust standard errors are clustered at the entrepreneur level and are shown in parentheses. p -values are shown in parentheses.

	Total Comments	Negative Comments	
Delay Duration	1.710 (0.002)	0.016 (0.000)	0.007 (0.197)
Project Experience			0.145 (0.846)
Delay Duration \times Project Experience			0.011 (0.022)
Failed Campaign Experience	55.540 (0.572)	1.479 (0.144)	1.896 (0.065)
Prior Campaign Funding Deviation	-4.176 (0.491)	-0.005 (0.933)	-0.004 (0.955)
Prior Project Delay	23.430 (0.319)	0.951 (0.008)	0.621 (0.067)
Execution Overlap	-377.799 (0.143)	-1.741 (0.201)	-3.133 (0.024)
External Financing	-518.389 (0.041)	4.748 (0.202)	4.202 (0.251)
New Category	-150.405 (0.191)	-1.322 (0.510)	-1.031 (0.614)
Elapsed Time	-0.632 (0.086)	0.006 (0.052)	0.003 (0.356)
Funding Period	-1.006 (0.834)	0.047 (0.382)	0.058 (0.279)
Funding Reward Tiers	-6.520 (0.365)	-0.190 (0.042)	-0.205 (0.022)
Funding Reward Size	-0.174 (0.048)	-0.004 (0.025)	-0.004 (0.026)
Ln Funding Threshold	-43.482 (0.415)	0.735 (0.184)	0.713 (0.191)
Ln Funding Exceeded	151.209 (0.046)	-0.039 (0.952)	-0.023 (0.971)
Ln Funding Backers	283.235 (0.015)	5.226 (0.000)	5.258 (0.000)
Baseline Updates	-12.661 (0.153)	0.087 (0.402)	0.080 (0.448)
Entrepreneur FE	Yes	Yes	Yes
Product Subtype FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Month FE	Yes	Yes	Yes
R ²	0.779	0.857	0.860
Entrepreneurs	303	303	303
Observations	644	644	644

Table A.12: Features and Negative Comments. Ordinary least squares (OLS) estimation. Robust standard errors are clustered at the entrepreneur level. *p*-values are shown in parentheses.

	Negative Comments					
Features Most	0.537 (0.441)			1.436 (0.071)		
Features Rank		0.597 (0.181)			1.236 (0.028)	
Features Percentile			0.722 (0.309)			1.729 (0.029)
Delay Duration				0.016 (0.019)	0.018 (0.062)	0.017 (0.015)
Features Most × Delay Duration				-0.002 (0.849)		
Features Rank × Delay Duration					-0.001 (0.743)	
Features Percentile × Delay Duration						-0.003 (0.706)
Project Experience	2.221 (0.059)	2.063 (0.078)	2.184 (0.063)	0.501 (0.585)	0.213 (0.812)	0.428 (0.640)
Failed Campaign Experience	1.522 (0.532)	1.461 (0.549)	1.504 (0.539)	-0.042 (0.970)	-0.285 (0.803)	-0.182 (0.873)
Prior Campaign Funding Deviation	-0.126 (0.115)	-0.126 (0.115)	-0.126 (0.116)	-0.106 (0.170)	-0.107 (0.164)	-0.106 (0.172)
Prior Project Delay	0.994 (0.103)	1.020 (0.095)	0.998 (0.103)	1.019 (0.036)	1.071 (0.027)	1.034 (0.033)
Execution Overlap	-5.064 (0.006)	-4.993 (0.006)	-5.016 (0.006)	-2.685 (0.101)	-2.623 (0.123)	-2.622 (0.104)
External Financing	7.285 (0.039)	7.313 (0.037)	7.219 (0.040)	5.528 (0.198)	5.645 (0.177)	5.596 (0.190)
New Category	-0.745 (0.730)	-0.873 (0.684)	-0.793 (0.712)	-1.232 (0.586)	-1.551 (0.485)	-1.346 (0.546)
Elapsed Time	0.005 (0.392)	0.005 (0.375)	0.005 (0.400)	0.008 (0.231)	0.008 (0.189)	0.008 (0.246)
Funding Period	0.098 (0.083)	0.096 (0.087)	0.096 (0.089)	0.089 (0.139)	0.086 (0.154)	0.087 (0.148)
Funding Reward Tiers	0.011 (0.913)	0.010 (0.919)	0.009 (0.931)	-0.052 (0.664)	-0.051 (0.671)	-0.054 (0.654)
Funding Reward Size	-0.006 (0.003)	-0.006 (0.003)	-0.006 (0.003)	-0.008 (0.001)	-0.008 (0.001)	-0.008 (0.001)
Ln Funding Threshold	2.067 (0.001)	2.052 (0.001)	2.069 (0.001)	1.878 (0.003)	1.875 (0.003)	1.892 (0.002)
Entrepreneur FE	Yes	Yes	Yes	Yes	Yes	Yes
Product Subtype FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
R ²	0.776	0.776	0.776	0.814	0.815	0.815
Entrepreneurs	314	314	314	303	303	303
Observations	722	722	722	644	644	644

image when people are saying you're delayed" (GPSCo CEO).

As an additional test for whether entrepreneurs continue to care about their image and the impact of negative comments (which are driven by delays), we look at whether entrepreneurs reduce the effort and care they put into customers over time. Specifically, we define a new variable, *Creator Engagement*, which is the total word count of the updates and comments posted by the entrepreneur on a given project. As shown in Table A.13, we find that there is a significant and positive relationship between *Creator Engagement* and *Project Experience*, indicating that entrepreneurs are posting and engaging with their backers more on each subsequent project. This suggests that there is no evidence in terms of their public engagement that entrepreneurs are less concerned with pleasing their customer base and defending their public reputation. In fact, the empirical evidence in our setting suggests the opposite. Again, the qualitative evidence aligns. TabletCo CEO described how even on his most recent project, he would "have people monitor the customer comments around the clock...to feed into consumer confidence."

A.1.12 INCENTIVE TO OVERPROMISE

Despite facing consequences for delay, it could be the case that the ex post consequences are offset by ex ante benefits. While making unrealistic predictions likely leads to negative consequences down the road when entrepreneurs fail to meet those predictions, if making aggressive predictions is helpful in the fundraising process entrepreneurs may still be incentivized to overpromise. We construct dependent variables for the amount of money raised in total and the amount of money raised in excess of the fundraising goal. We regress these dependent variables on an independent variable of the predicted delivery time. If there was a benefit to overpromising, we would expect to see a negative relationship here, with an increase in promised delivery time decreasing the amount of funds raised. As shown in Table B.7, we do not find a statistically significant relationship, suggesting that there is no general

Table A.13: Creator Engagement. Ordinary least squares (OLS) estimation. Robust standard errors are clustered at the entrepreneur level. *p*-values are shown in parentheses.

	Creator Engagement
Project Experience	1660.434 (0.021)
Failed Campaign Experience	1510.288 (0.019)
Prior Campaign Funding Deviation	9.544 (0.653)
Prior Project Delay	-462.340 (0.023)
Execution Overlap	-1931.397 (0.237)
External Financing	687.536 (0.508)
New Category	719.917 (0.219)
Elapsed Time	-2.272 (0.241)
Funding Period	4.890 (0.820)
Funding Reward Tiers	-79.848 (0.072)
Funding Reward Size	-1.240 (0.067)
Ln Funding Threshold	1247.513 (0.071)
Ln Funding Exceeded	1187.129 (0.011)
Ln Funding Backers	-836.528 (0.284)
Baseline Updates	215.579 (0.034)
Entrepreneur FE	Yes
Product Subtype FE	Yes
Year FE	Yes
Month FE	Yes
R ²	0.720
Entrepreneurs	306
Observations	626

incentive to make unrealistic predictions for the sake of upfront financing.⁷ This is of course peculiar given the assumed desire of customers to get their products faster, but customers may mentally discount aggressive delivery times because they know they are unrealistic and thus not place any value on them. Overall, we contend that overpromising aggressive delivery times is not the major driver of the main findings.

A.1.1.3 EXITING AFTER VC FINANCING

We then turn to whether certain types of entrepreneurs are exiting our sample, looking first at whether entrepreneurs are more likely to exit the sample after raising external capital from another source (generally an indicator of higher quality or ability). We test whether high-quality entrepreneurs might exit the sample. In particular, we test the assumption that these higher-quality entrepreneurs may desire and be able to raise venture capital financing and leave crowdfunding. We define a binary indicator of whether the entrepreneur goes on to do a subsequent project as the dependent variable. For projects completed towards the end of our sample timeline, it is unclear whether the entrepreneur has truly exited the sample or is in the process of preparing another campaign. As such, we drop the projects from the last year of our sample when performing this analysis. We then regress the binary indicator of going on to do a subsequent project on a binary indicator of whether the entrepreneur has raised external capital at that time. As shown in Table B.9, we find that raising venture funding does not impact whether entrepreneurs exit the sample or continue on and do another Kickstarter project.

⁷We recognize that the product specification is likely correlated with both the predicted time and the funds raised (e.g., really fancy product specifications have longer predicted times, but also raise more money). We include *Features Most* in the regression to try to serve as a proxy for the product specification; however, given that this is an imprecise measure of product specification there is likely still some omitted variable bias. The results are similar if we include *Features Ranks* or *Features Percentile* as controls.

Table A.14: Incentive to Overpromise. Ordinary least squares (OLS) estimation. Robust standard errors are clustered at the entrepreneur level. *p*-values are shown in parentheses.

	Ln Funds Raised	Ln Funding Exceeded
Predicted Time	0.001 (0.102)	0.002 (0.247)
Features Rank	0.111 (0.026)	0.169 (0.053)
Failed Campaign Experience	-0.097 (0.319)	-0.196 (0.149)
Prior Campaign Funding Deviation	-0.022 (0.001)	-0.026 (0.000)
Prior Project Delay	0.024 (0.454)	0.044 (0.435)
Execution Overlap	-0.163 (0.212)	-0.349 (0.125)
External Financing	0.137 (0.680)	0.302 (0.660)
New Category	0.023 (0.906)	-0.172 (0.559)
Project Experience	0.082 (0.182)	0.044 (0.665)
Elapsed Time	0.001 (0.509)	0.001 (0.477)
Funding Period	0.012 (0.002)	0.014 (0.005)
Funding Reward Tiers	0.023 (0.004)	0.031 (0.025)
Funding Reward Size	0.000 (0.198)	0.000 (0.328)
Ln Funding Threshold	0.405 (0.000)	0.015 (0.843)
Entrepreneur FE	Yes	Yes
Product Subtype FE	Yes	Yes
Year FE	Yes	Yes
Month FE	Yes	Yes
R ²	0.923	0.840
Entrepreneurs	314	314
Observations	722	722

Table A.15: Exiting After VC Financing. Ordinary least squares (OLS) estimation. Robust standard errors are clustered at the entrepreneur level. p -values are shown in parentheses.

	Does Another Project
External Financing	0.033 (0.819)
Failed Campaign Experience	0.052 (0.617)
Prior Campaign Funding Deviation	-0.004 (0.124)
Prior Project Delay	-0.028 (0.132)
Execution Overlap	-0.081 (0.408)
New Category	-0.199 (0.042)
Project Experience	-0.394 (0.000)
Elapsed Time	-0.001 (0.001)
Funding Period	-0.003 (0.237)
Funding Reward Tiers	0.008 (0.195)
Funding Reward Size	0.000 (0.217)
Ln Funding Threshold	-0.035 (0.254)
Ln Funding Exceeded	0.030 (0.354)
Ln Funding Backers	-0.005 (0.921)
Entrepreneur FE	Yes
Product Subtype FE	Yes
Year FE	Yes
Month FE	Yes
R ²	0.770
Entrepreneurs	304
Observations	641


A.1.14 EXITING AFTER DELAY

We now evaluate whether entrepreneurs disproportionately exit the sample after they delay, which might suggest that entrepreneurs who learn leave while those who do not learn stay. Using the same dependent variable of whether the entrepreneur goes on to do another project and the same sample as the prior section, we regress this indicator of going on to do another project on the delay duration of the current project and do not find a significant effect. This result is shown in Table A.16. As an additional check, we also looked at the distribution of the difference in delay between the second project of entrepreneurs who went on to do a third project matched with the most similar second project by an entrepreneur who did not go on to do a third project. If entrepreneurs left the sample after they delayed, we would expect to see a distribution skewed toward negative values (implying that the second project delay of those who went on to do a third project is smaller than the second project delay of those who exited the sample after two projects). However, we see a roughly normal distribution centered around zero.

Table A.16: Exiting After Delay. Ordinary least squares (OLS) estimation. Robust standard errors are clustered at the entrepreneur level. p -values are shown in parentheses.


	Does Another Project
Delay Duration	-0.0003 (0.194)
Failed Campaign Experience	0.235 (0.116)
Prior Campaign Funding Deviation	-0.004 (0.159)
Prior Project Delay	-0.024 (0.337)
Execution Overlap	-0.034 (0.785)
New Category	-0.266 (0.049)
Project Experience	-0.344 (0.000)
Elapsed Time	-0.001 (0.016)
Funding Period	-0.003 (0.352)
Funding Reward Tiers	0.008 (0.238)
Funding Reward Size	0.000 (0.527)
Ln Funding Threshold	-0.042 (0.259)
Ln Funding Exceeded	0.045 (0.280)
Ln Funding Backers	-0.014 (0.821)
Entrepreneur FE	Yes
Product Subtype FE	Yes
Year FE	Yes
Month FE	Yes
R ²	0.780
Entrepreneurs	294
Observations	581

Ringo Robot



Smart Parts

- 1 Accelerometer
- 2 Motor Driver
- 3 Battery Charger
- 4 Gyroscope
- 5 Arduino UNO MCU



Ins and Outs

- 6 Ambient Light Sensor
- 7 RGB NeoPixel LED
- 8 IR LED Transmitter
- 9 38kHz TV Remote Receiver
- 10 Piezo Sound Element
- 11 Power Switch
- 12 Reset Button
- 13 User Button
- 14 USB Port for Charging
- 15 Charging Status LED
- 16 Programming Port
- 17 Edge/Line Sensor
- 18 IR LED Surface Illuminator



Spirit Rover








Figure A.1: Plum Geek Robotics Product Features. Reproduction of breakdowns of product features provided by Plum Geek Robotics on its company website. Ringo Robot was its first Kickstarter project, launched on February 25, 2015. The Spirit Rover was its third project, launched on September 28, 2016.


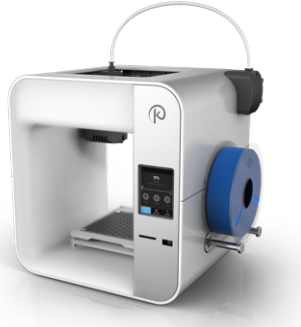
Low Complexity	High Complexity
	
Components	Components
<ul style="list-style-type: none"> Metal casing Left bevel Wrench cutout Written measurements 	<ul style="list-style-type: none"> Nozzle Printing bed Heating apparatus Camera LED lights Color touchscreen Bearing spool holder Casing Power cords Circuit board Software operating system
Interdependencies to be Addressed	Interdependencies to be Addressed
<ul style="list-style-type: none"> Metal casing dimensions match specified measurements. Metal casing will fit in the machine that adds the measurements. Writing color will show up and adhere to the metal surface. Separate machine can accommodate metal casing to add left bevel. Separate machine can accommodate metal casing to add wrench hole. Intended functions of the metal casing are not impaired by cutout wrench hole. Upper ruler cutout does not fall off or weaken the tool's structural integrity. 	<ul style="list-style-type: none"> <i>Only subset included for heating apparatus, bearing spool holder, and casing.</i> Power supply is sufficient for heating apparatus, camera, LEDs, touch screen, etc. Bearing spool holder can hold weight of spool and remains balanced with added weight. Heating mechanism does not impair function of screen or other electronics. Heating apparatus creates correct temperature for input material. Casing accommodates LED lights, printing bed, nozzle, touch screen, etc. Casing provides appropriate spacing between camera and printing bed for optimal focus. Casing allows nozzle sufficient range of motion when printing. Bearing spool holder fits the input material. Bearing spool holder attaches to casing.

Figure A.2: Project Complexity Examples. Selected projects from sample to illustrate a typical “low” complexity and “high” complexity project. *Components* (or modules) and *Interdependencies to be Addressed* represent only those confirmed by the researchers and likely constitute only a subset of those actually faced by the entrepreneurs. In the context of the NK model, the separate *components* represent N, and *Interdependencies to be Addressed* represent (an observable subset of) K.

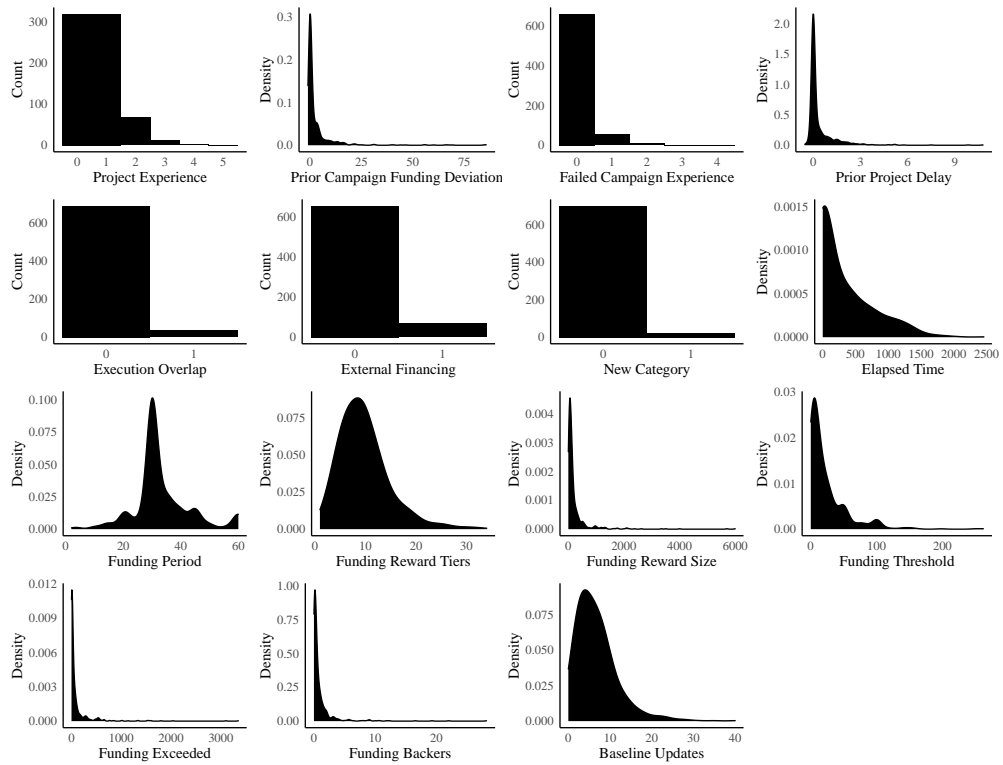


Figure A.3: Variable Distributions. Visualizations of the distributions of independent variables, with density plots for continuous measures and histograms for those with finite discrete values.

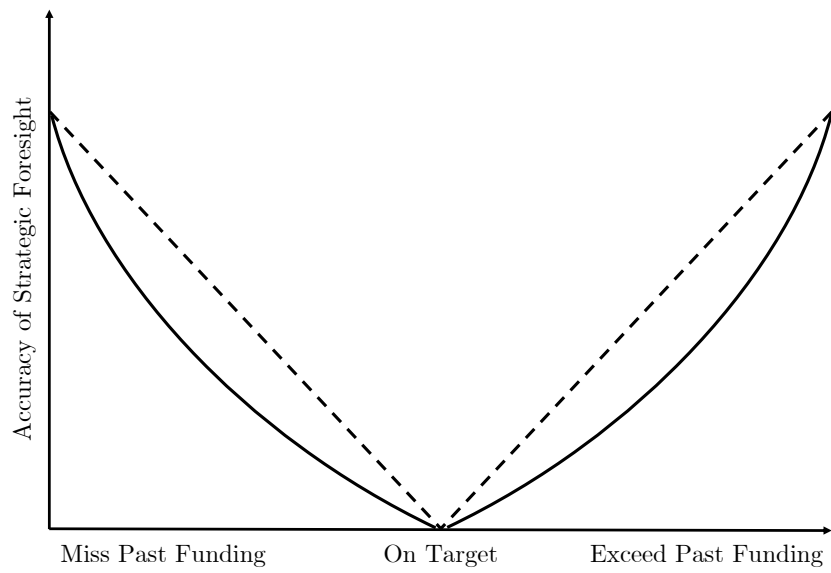


Figure A.4: Outcome–Aspiration Gap and Accuracy of Strategic Foresight. In our context, we operationalize the accuracy of strategic foresight as *Unforeseen Interdependencies* and *Delay Duration*.




Figure A.5: Deviation and Delay. The binned scatter and linear fit to the left of 0% *Prior Campaign Funding Deviation* show the relationship with *Delay Percent* on the current project and the extent to which the entrepreneur failed to meet her funding threshold on the prior project. To the right, they show the same relationship but for the extent to which the entrepreneur met or exceeded her funding threshold on the prior project.




Figure A.6: Deviation and Unforeseen Interdependencies. The binned scatter and linear fit to the left of 0% *Prior Campaign Funding Deviation* show the relationship with *Unforeseen Interdependencies* on the current project and the extent to which the entrepreneur failed to meet her funding threshold on the prior project. To the right, they show the same relationship but for the extent to which the entrepreneur met or exceeded her funding threshold on the prior project.

B

Complexity and Timing in Novel Ventures

B.1 APPENDIX

B.1.1 VALIDATING MEASURE OF NOVEL RECOMBINATION

Our measure of novel recombination is derived using the keywords and descriptions of Pitchbook firms. These keywords and descriptions are created through the joint efforts of the entrepreneurial firms, venture investors, the broader startup community, and Pitchbook employees. Despite the many

different parties involved in capturing this data, there is of course room for human error or inconsistencies. As such, it is essential to validate our measure of novelty. By comparing our measure to other proxies of novelty, namely entry order and tech stack novelty, we find strong correlations that help to validate our measure.

ENTRY ORDER IN COMPETITIVE CLUSTERS

We first use entry order within competitive clusters to validate our measure of novelty. We use competitive clusters as a proxy for market opportunities; identifying the set of firms competing against each other allows us to partition the firms into different clusters, which proxy different opportunities. If our measure of novelty is valid, we would assume that the first firms to compete for a given opportunity would be more likely to have novel recombinations.

Each firm in our sample lists a set of competitors. We use this data to create a directed graph, with each firm pointing to its identified competitors. From here, there are many different approaches to break this directed graph into clusters. These different approaches balance considerations such as whether to allow a company to appear in multiple clusters, maximum cluster size, maximum computing time, memory requirements, etc. Given the various trade-offs across these different dimensions, we employ four different techniques to create clusters and compare each technique to our novelty measure. These techniques include:

WALKTRAP ALGORITHM (Pons & Latapy, 2006) The Walktrap Algorithm uses random walks to identify clusters. The basic intuition of the algorithm is that random walks on a graph tend to get trapped into densely connected parts, which should correspond to clusters. The length of the random walk is an input parameter. After conducting multiple random walks, a similarity matrix is constructed. This matrix measures the similarity between pairs of nodes based on how often they appear together during the random walks. The matrix is then used to construct a hierarchical tree-like struc-

ture (a dendrogram) that represents the network's cluster structure, which is cut at a certain level to determine the number of clusters in the network.

CLIQUE METHOD (Fortunato & Barthelemy, 2007) The Clique Method identifies cliques, which are groups of nodes where each node is directly connected to every other node in the group. In other words, this method regards each maximal clique (complete subgraph that is not a subgraph of another complete graph) as a cluster. According to this definition, each company in a cluster is connected to any other company in the cluster.

LEIDEN ALGORITHM (Traag et al., 2019) The Leiden Algorithm is often used for single-cell RNA sequencing data analysis. It operates by optimizing a quality function that measures the significance of node assignments to clusters. Leiden starts with an initial clustering and iteratively evaluates the quality of the clusters by considering the benefits of moving nodes between clusters. If moving nodes enhances cluster quality, the algorithm makes that adjustment. This iterative process continues until no further improvements can be made.

INFOMAP ALGORITHM (Rosvall & Bergstrom, 2008) The Infomap Algorithm analyzes the network's structure and seeks to minimize the information needed to describe it. It starts by exploring the network using random walks, tracking how information flows between nodes. The algorithm identifies clusters by finding regions where information can be efficiently compressed. It optimizes these clusters by iteratively merging or splitting them to minimize the description length.

The performance and output of these different clustering methods are summarized in Table C.1. We see that the Walktrap Algorithm generates very big clusters. Using the default length of the random walk (with a value of 4), the biggest cluster has a size of 3,009. In our context, imagining an opportunity space of over 3,000 firms seems very large. Increasing the walk length to improve the results makes it computationally too slow.

Table B.1: Clustering Method Performance Summary. Summary statistics for the four different clustering methods used to determine the different sets of companies pursuing a given opportunity.

	<i>Clique</i>	<i>Walktrap</i>	<i>Infomap</i>	<i>Leiden</i>
Number of Clusters	70,561	6,360	5,574	4,101
Minimum Cluster Size	2	2	2	2
Maximum Cluster Size	16	3,009	414	196
Mean Cluster Size	4.071	7.088	8.088	10.992
Median Cluster Size	3	3	3	2
Number of Clusters of Size ≥ 2	22,015	3,166	2,586	2,396
Number of Clusters of Size ≥ 100	0	36	26	97

On the other hand, the Clique Method generates many small clusters. This is because a company can be classified into numerous cliques. A strength of this method is that it is straightforward and does a good job of putting together groups of reasonable size such that all the firms would be aware of each other, though there is certainly some overlap between clusters. Leiden is scalable, making it suitable for analyzing large networks. We adjusted the resolution parameter to 28, which lowered the maximum cluster size to a more reasonable number. The Infomap Algorithm is also computationally efficient: it is fast and generates reasonable cluster sizes.

For each of the clustering methods, we compute measures to capture when each company started to pursue the defined opportunity. *Entry Order* is the sequential entry order of companies in the cluster pursuing the opportunity. *Years Since First Entrant* is the time in years since the first firm in the cluster was founded to pursue the opportunity. *Enter in First Year* is a binary indicator variable equal to 1 if the firm started pursuing a given opportunity in the first year, and zero otherwise. *Enter in First Two Years* is a binary indicator variable equal to 1 if the firm started pursuing a given opportunity in the first two years, and zero otherwise.

Table C.9 outlines the summary statistics of these measures for each of the clustering methods. Table C.10 shows a correlation matrix. As might be expected, we see that the entry order variables using different clustering methods are generally correlated—this seems to validate that these measures

are capturing a similar concept, but with some distinctions given the priorities of the different clustering algorithms.

Table B.2: Summary Statistics by Clustering Method. 27,847 firm-level observations.

Statistic	Mean	St. Dev.	Min	Max
Entry Order (Walktrap)	190.61	485.21	1	2,950
Years Since First Entrant (Walktrap)	52.42	58.81	0	207
Enter in First Year (Walktrap)	0.18	0.38	0	1
Enter in First Two Years (Walktrap)	0.20	0.40	0	1
Entry Order (Clique)	2.24	1.41	1.00	14.00
Years Since First Entrant (Clique)	8.78	17.58	0	199
Enter in First Year (Clique)	0.43	0.50	0	1
Enter in First Two Years (Clique)	0.49	0.50	0	1
Entry Order (Leiden)	34.01	33.60	1	192
Years Since First Entrant (Leiden)	48.70	49.98	0	212
Enter in First Year (Leiden)	0.11	0.32	0	1
Enter in First Two Years (Leiden)	0.13	0.33	0	1
Entry Order (Infomap)	20.66	33.62	1	399.50
Years Since First Entrant (Infomap)	32.41	39.73	0	207
Enter in First Year (Infomap)	0.16	0.37	0	1
Enter in First Two Years (Infomap)	0.18	0.38	0	1

Table B.3: Pairwise Correlation Matrix by Clustering Method.

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
Entry Order (Walktrap)	1															
Years Since First Entrant (Walktrap)	0.638	1														
Enter in First Year (Walktrap)	-0.182	-0.414	1													
Enter in First Two Years (Walktrap)	-0.196	-0.444	0.928	1												
Entry Order (Clique)	0.257	0.275	-0.341	-0.340	1											
Years Since First Entrant (Clique)	-0.014	0.221	-0.210	-0.223	0.269	1										
Enter in First Year (Clique)	-0.160	-0.219	0.488	0.463	-0.607	-0.426	1									
Enter in First Two Years (Clique)	-0.148	-0.222	0.437	0.468	-0.581	-0.466	0.891	1								
Entry Order (Leiden)	0.265	0.324	-0.315	-0.321	0.390	0.084	-0.315	-0.295	1							
Years Since First Entrant (Leiden)	-0.044	0.317	-0.257	-0.267	0.146	0.309	-0.182	-0.195	0.421	1						
Enter in First Year (Leiden)	-0.138	-0.299	0.736	0.682	-0.277	-0.175	0.409	0.365	-0.349	-0.347	1					
Enter in First Two Years (Leiden)	-0.145	-0.313	0.695	0.720	-0.276	-0.183	0.386	0.386	-0.368	-0.366	0.946	1				
Entry Order (Infomap)	0.401	0.434	-0.247	-0.262	0.367	0.051	-0.235	-0.224	0.541	0.145	-0.206	-0.217	1			
Years Since First Entrant (Infomap)	0.049	0.483	-0.305	-0.327	0.200	0.409	-0.229	-0.247	0.296	0.662	-0.289	-0.303	0.291	1		
Enter in First Year (Infomap)	-0.153	-0.322	0.815	0.753	-0.325	-0.201	0.468	0.418	-0.329	-0.293	0.817	0.773	-0.250	-0.352	1	
Enter in First Two Years (Infomap)	-0.163	-0.342	0.760	0.805	-0.326	-0.214	0.449	0.447	-0.336	-0.304	0.758	0.801	-0.267	-0.377	0.928	1

We then regress our measure of *Novel Recombination* on *Entry Order*, *Years Since First Entrant*, *Enter in First Year*, and *Enter in First Two Years* as defined under each of the four different clustering methods. As reported in Table C.3, we find a consistent and strong correlation between our measure of novelty and each of these measures. Using the Clique Method, we find that a 1% increase in *Entry*

Order leads to a 0.541% decrease in *Novel Recombination* ($p < .001$), a 1% increase in *Years Since First Entrant* leads to a 0.171% decrease in *Novel Recombination* ($p < .001$), entering in the first year leads to a 32.8% increase in *Novel Recombination* ($p < .001$), and entering in the first two years leads to a 36.4% increase in *Novel Recombination* ($p < .001$). Using the Walttrap Algorithm, we find that a 1% increase in *Entry Order* leads to a 0.359% decrease in *Novel Recombination* ($p < .001$), a 1% increase in *Years Since First Entrant* leads to a 0.253% decrease in *Novel Recombination* ($p < .001$), entering in the first year leads to a 47.9% increase in *Novel Recombination* ($p < .001$), and entering in the first two years leads to a 48.8% increase in *Novel Recombination* ($p < .001$). Using the Leiden Method, we find that a 1% increase in *Entry Order* leads to a 0.374% decrease in *Novel Recombination* ($p < .001$), a 1% increase in *Years Since First Entrant* leads to a 0.312% decrease in *Novel Recombination* ($p < .001$), entering in the first year leads to a 57.9% increase in *Novel Recombination* ($p < .001$), and entering in the first two years leads to a 59.6% increase in *Novel Recombination* ($p < .001$). Using the Infomap Method, we find that a 1% increase in *Entry Order* leads to a 0.386% decrease in *Novel Recombination* ($p < .001$), a 1% increase in *Years Since First Entrant* leads to a 0.268% decrease in *Novel Recombination* ($p < .001$), entering in the first year leads to a 52.0% increase in *Novel Recombination* ($p < .001$), and entering in the first two years leads to a 52.3% increase in *Novel Recombination* ($p < .001$).

The strong and consistent relationship we find between our measure of novelty and those derived using entry order within opportunity clusters provides compelling support that our measure is indeed capturing the intended mechanism.

TECH STACK INNOVATION

As another means of validating our measure of novel recombination, we compare our measure to the frequency of novel combinations in the product tech stack. We assume that firms implementing an opportunity with more novel recombinations will use more novel combinations of APIs in their

Table B.4: Novel Recombination and Entry Order. Ordinary least squares (OLS) estimation. Robust standard errors are clustered at the clique or cluster level. p -values are shown in brackets.

	Ln Novel Recombination	Ln Novel Recombination	Ln Novel Recombination	Ln Novel Recombination
Panel A: Clique Method				
Ln Entry Order	-0.541 [<0.001]			
Ln Years Since First Entrant		-0.171 [<0.001]		
Enter in First Year (Binary)			0.328 [<0.001]	
Enter in First Two Years (Binary)				0.364 [<0.001]
FE: Clique	X	X	X	X
Observations	27 847	27 847	27 847	27 847
R^2	0.873	0.875	0.871	0.872
Panel B: Walktrap Method				
Ln Entry Order	-0.359 [<0.001]			
Ln Years Since First Entrant		-0.253 [<0.001]		
Enter in First Year (Binary)			0.479 [<0.001]	
Enter in First Two Years (Binary)				0.488 [<0.001]
FE: Walktrap Cluster	X	X	X	X
Observations	27 847	27 847	27 847	27 847
R^2	0.355	0.336	0.312	0.313
Panel C: Leiden Method				
Ln Entry Order	-0.374 [<0.001]			
Ln Years Since First Entrant		-0.312 [<0.001]		
Enter in First Year (Binary)			0.579 [<0.001]	
Enter in First Two Years (Binary)				0.596 [<0.001]
FE: Leiden Cluster	X	X	X	X
Observations	27 847	27 847	27 847	27 847
R^2	0.291	0.264	0.230	0.230
Panel D: Infomap Method				
Ln Entry Order	-0.386 [<0.001]			
Ln Years Since First Entrant		-0.268 [<0.001]		
Enter in First Year (Binary)			0.520 [<0.001]	
Enter in First Two Years (Binary)				0.523 [<0.001]
FE: Infomap Cluster	X	X	X	X
Observations	27 847	27 847	27 847	27 847
R^2	0.336	0.325	0.295	0.296

tech stack; in the process of realizing more novel recombinations, these firms will be more likely to bring together components that other firms have not used together historically. This validation is of course only relevant to the subset of companies where the product is a website, such that the tech stack reflects the firm's efforts to realize the novel recombination. In particular, we look at the subset of firms from our sample who have some API offering, making the assumption that firms that offer an API integration are offering web-based services such that their website would be indicative of their product.

To construct our measures, we compute all the combinations of each firm's tech stack components along with the date on which the firm first implemented that combination. We define *Total Novel Tech Combinations* as the count of how many of a firm's tech stack combinations are novel, meaning the firm was the first firm in the sample to use the combination. As a more continuous measure, we can also look at how many other firms had implemented a given tech combination previously as a proxy for how dispersed the combination was at the time the focal firm implemented it. *Tech Combinations Dispersion (Mean)* takes the mean (across all of a firm's tech stack component combinations) number of times other firms had already implemented each combination. *Tech Combinations Dispersion (Median)* is the same measure, but instead uses the median across each firm's set of component combinations.

Table C.2 outlines both the summary statistics and correlations. We see that *Total Novel Tech Combinations* has a very wide spread and large standard deviation. This is related to how total combinations between components explodes exponentially as the total components increase. As such, we take the natural log of this variable in any regressions. Further, we see a meaningful negative correlation between the measure of novelty and the two measures of dispersion. This is expected; an API combination this common and widely dispersed is of course not going to be novel. And, as is expected, the mean and median measures of dispersion are also highly correlated.

We then regress both our binary and continuous measures of novel recombination on each of

Table B.5: Tech Stack Summary Statistics and Pairwise Correlation.

	Summary Stats				Correlations		
	Mean	St. Dev.	Min	Max	Total Novel Tech Combinations	Tech Combination Dispersion (Mean)	Tech Combination Dispersion (Median)
Total Novel Tech Combinations	3,759.25	4,942.45	0	93,152	1		
Tech Combination Dispersion (Mean)	24.32	7.83	1	59.08	-0.463	1	
Tech Combination Dispersion (Median)	15.10	8.19	1	62	-0.413	0.936	1

these measures of novel tech stack combinations. Interpreting the logged variables presented in Table B.6 in words, every doubling of the number of novel API combinations increases the likelihood of having a novel recombination by 2.4% ($p = 0.004$) and increases the total novel recombinations by 3.5% ($p = 0.012$). Increasing a firm’s mean dispersion across its API combination by one decreases the likelihood that the firm has any novel recombinations by 0.4 percentage points ($p = 0.014$); across the full range from the minimum to maximum value (from 1 to 59) in the sample this creates a swing of 25 percentage points. This same increase also leads to a 0.4% ($p = 0.159$) decrease in the total number of novel recombinations. Increasing a firm’s median dispersion across its API combination by one decreases the likelihood that the firm has any novel recombinations by 0.5 percentage points ($p = 0.005$); across the full range from the minimum to maximum value (from 1 to 62) in the sample this creates a swing of 31 percentage points. This same increase also leads to a 0.5% ($p = 0.132$) decrease in the total number of novel recombinations.

Table B.6: Novel Recombination and Tech Stack Innovation. Ordinary least squares (OLS) estimation with robust standard errors. p -values are shown in brackets.

	Novel Recombination Binary	Novel Recombination Binary	Novel Recombination Binary	Ln Novel Recombination	Ln Novel Recombination	Ln Novel Recombination
Ln Total Novel Tech Combinations	0.035 [0.004]			0.049 [0.012]		
Tech Combination Dispersion (Mean)		-0.004 [0.014]			-0.004 [0.159]	
Tech Combination Dispersion (Median)			-0.005 [0.005]			-0.005 [0.132]
Observations	1378	1378	1378	1378	1378	1378
R ²	0.006	0.004	0.006	0.004	0.001	0.002

As noted previously in the paper, we find that firms with more novel recombination, on average, use fewer APIs: likely because their products are less modular and more things need to be built from scratch. However, the analysis here shows that, while the firms do use fewer APIs, on average the combinations of APIs they use are more likely to be novel. This correlation between novel recombination (generated from combinations of words in the company’s description) and novel combinations of APIs suggests that the company implements more novel interdependencies when implementing opportunities containing novel recombinations.

B.1.2 LINKEDIN SAMPLE CONSTRUCTION

The LinkedIn data starts at the job level and contains detailed information on the employee, start date, end date, job description, firm, etc. We eliminate all employee jobs where there is no start date or end date. For jobs with only a start year, we set the end year equal to the start year. For jobs with only an end year, we set the start year equal to the end year. We also eliminate any jobs where the start date comes after the end date or employees with job start dates more than two years prior to the firm’s founding year. Consistent with our performance analysis, we only look at the first ten funding rounds (given that going beyond ten rounds generally indicates abnormal fundraising behavior—a “Series K” round is not common). For jobs that are listed as the employee’s current job, we set the end date to December 2020 (the date the data was pulled).

A little over a third of jobs include a job description. We look at the words in each job description and count the number of matches against the set of keywords associated with managing complexity and managing third parties (detailed in Appendix Section B.1.3). These counts are then aggregated to the firm-year-level. We use the firm website URL listed in Pitchbook as the key to match firms with the LinkedIn data. We find a match for 16,892 firms (54% of the total sample of 31,450 firms).

B.1.3 *MANAGE COMPLEXITY AND MANAGE THIRD PARTIES* VARIABLE CONSTRUCTION

To construct our proxy of complexity, we search the job descriptions of employees working at the firms in our sample for words that suggest managing interdependencies and working with complexity. We determined the set of relevant words by aggregating all the words in job descriptions and looking for the most frequent words. We iterated between the possible words and the set of words the extant literature generally uses when discussing complexity.

The set of terms includes the following words: across, between, combination, combinations, combine, combined, combines, combining, conjoint, connect, connecting, connects, coordinate, coordinated, coordinates, coordinating, depend, dependencies, dependency, dependent, integrate, integrated, integrates, integrating, interactive, interchange, interchanged, interchanges, interchanging, interconnect, interconnected, interconnecting, interconnects, interdepartment, interdepend, interdependent, interdepending, interdisciplinary, intermediaries, intermediary, intermediate, intermediates, intermesh, intermeshed, intermolecular, interop, interoperable, interplay, interplays, interrelate, interrelated, interrelating, intersec, intersect, intersecting, intersection, intersections, intersects, joint, jointly.

In addition to looking at complexity, we use the same process (looking at words by frequency along with the extant literature) to identify markers of managing third parties. These words include: alliance, alliances, alliancing, buyer, buyers, complementors, complements, contractor, contractors, customer, customers, external, government, intercity, intercompanies, intercompany, intercultural, international, interstate, lobbyist, network, networks, partner, partners, partnership, partnerships, platform, platforms, regulator, regulators, relations, seller, sellers, supplier, suppliers, supplies, supply.

B.1.4 CORRELATION TABLES

Table B.7 contains correlation statistics for the variables at the firm-year level. First, we observe a high correlation between *Manage Complexity* and *Manage Complexity (Count)* (0.957) as well as between *Manage Third Parties* and *Manage Third Parties (Count)* (0.970). Because the variables within each pair of terms are intended to capture the same effect, this high correlation makes sense and validates that they do in fact capture a similar effect. We also see strong correlations between the two paring groups. Again, this is to be expected given that managing third parties is a form of managing complexity. *Round* and *Firm Age* are correlated (0.467) given that older firms tend to be raising higher numbered rounds. We also find a correlation between *Novel Recombination* and *Firm Scope* (0.619). This makes sense given that firms with a larger scope are more likely to have a larger number of novel combinations. We also observe correlations between *Total Employees*, *Total Descriptions*, and *Description Words*. Again this is by construction as more employees at a firm will, on average, increase the number of employees with job descriptions. This will, in turn and on average, increase the total number of words across all descriptions.

Table B.7: Firm-Year Panel Pairwise Correlation Matrix.

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
(1) Manage Complexity	1														
(2) Manage Complexity (Count)	0.957	1													
(3) Manage Third Parties	0.838	0.864	1												
(4) Manage Third Parties (Count)	0.862	0.915	0.970	1											
(5) Novel Recombination	-0.029	-0.030	-0.023	-0.026	1										
(6) Firm Age	0.127	0.132	0.102	0.117	0.077	1									
(7) Firm Scope	-0.030	-0.032	-0.024	-0.028	0.619	-0.077	1								
(8) Cumulative Funding	0.116	0.129	0.132	0.139	-0.010	0.056	-0.002	1							
(9) Round Size	0.052	0.059	0.051	0.056	-0.007	0.013	-0.003	0.333	1						
(10) CEO Change	-0.0001	0.0004	-0.001	-0.0003	0.009	-0.030	0.009	0.010	0.023	1					
(11) Total Employees	0.679	0.743	0.731	0.778	-0.022	0.133	-0.015	0.181	0.068	-0.005	1				
(12) Total Descriptions	0.839	0.916	0.908	0.957	-0.025	0.126	-0.031	0.150	0.063	-0.001	0.842	1			
(13) Description Words	0.857	0.918	0.939	0.973	-0.026	0.132	-0.034	0.138	0.060	-0.001	0.786	0.974	1		
(14) Year	-0.003	-0.006	-0.008	-0.011	-0.029	0.245	0.132	0.029	0.011	-0.018	0.030	-0.018	-0.015	1	
(15) Round	0.139	0.151	0.104	0.125	-0.001	0.467	-0.042	0.136	0.052	0.093	0.154	0.138	0.142	0.235	1

Table B.8 contains correlation statistics for the variables at the firm-round level. First, we observe a high correlation between *Need Bridge Financing* and both *Receive Bridge Financing* (0.813) and *Shut Down* (0.535). This is to be expected given that all firms that receive a bridge or fail both needed

a bridge. We also find a correlation between *Novel Recombination* and *Firm Scope* (0.597). This makes sense given that firms with a larger scope are more likely to have a larger number of novel combinations. We also observe a high correlation (0.571) between *Round* and *Firm Age*. Given that these variables are mechanically related with more time elapsing since founding with each subsequent funding round, this is expected. We also observe a high correlation between *Round Size*, *Cumulative Funding Amount*, *Premoney Valuation*, and *Postmoney Valuation*. Again, this is expected given the mechanical relationship between these variables (e.g., the larger the round size, the larger the postmoney valuation is by construction, with firms tending to raise larger amounts in later rounds).

Table B.8: Firm-Round Panel Pairwise Correlation Matrix.

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
(1) Receive Bridge Financing	1															
(2) Need Bridge Financing	0.813	1														
(3) Shutdown	-0.039	0.535	1													
(4) IPO	-0.032	-0.036	-0.016	1												
(5) Premoney Valuation	-0.008	-0.011	-0.008	0.049	1											
(6) Postmoney Valuation	-0.008	-0.012	-0.008	0.051	0.989	1										
(7) Novel Recombination	0.008	0.001	-0.011	-0.007	-0.010	-0.011	1									
(8) Novel Recombination Binary	0.011	0.010	0.001	-0.002	-0.005	-0.006	0.416	1								
(9) Firm Scope	-0.035	-0.061	-0.056	-0.019	0.018	0.017	0.597	0.299	1							
(10) Round Size	-0.008	-0.010	-0.004	0.042	0.455	0.577	-0.012	-0.014	0.004	1						
(11) Cumulative Funding	-0.005	-0.008	-0.006	0.036	0.405	0.465	-0.012	-0.017	0.006	0.552	1					
(12) Firm Age	0.065	0.082	0.047	0.106	0.064	0.066	0.099	0.123	0.004	0.055	0.059	1				
(13) CEO Change	0.032	0.038	0.021	0.008	-0.004	-0.004	0.028	0.025	0.006	0.004	0.009	0.148	1			
(14) Round	0.052	0.063	0.037	0.095	0.128	0.133	0.032	0.053	0.029	0.097	0.140	0.571	0.169	1		
(15) Round Year	-0.077	-0.077	-0.017	0.016	0.046	0.047	-0.024	-0.101	0.190	0.028	0.034	0.138	-0.003	0.191	1	
(16) Founding Year	-0.109	-0.120	-0.045	-0.053	-0.00001	-0.001	-0.084	-0.166	0.162	-0.011	-0.008	-0.516	-0.096	-0.198	0.776	1

B.1.5 VALIDATING COMPLEXITY AND ITS RELATIONSHIP WITH NOVEL RECOMBINATION

For the subset of firms who produce a software/Software-as-a-Service (SaaS) product, we also collect data from BuiltWith on the application programming interfaces (APIs) used to build their product. For this subset of 6,208 firms, we collect a monthly panel of which APIs each firm uses on its website (in the case of the SaaS companies we subset to, the website is the product), resulting in a panel of 760,678 observations at the firm-month level with an average of 80 month-level observations for each firm. Summary statistics and correlations appear in Table B.9.

Table B.9: Software Firm-Month Panel Summary Statistics and Pairwise Correlation. 760,678 firm-month-level observations of 6,208 unique firms with industry or vertical labels as software or SaaS.

	Summary Stats				Correlations		
	Mean	St. Dev.	Min	Max	API Count	Months	Novel Recombination (Binary)
API Count	56.59	77.38	0	4,203	1		
Months	80.20	60.97	1	282	0.337	1	
Novel Recombination (Binary)	0.56	0.50	0	1	0.010	0.030	1

APIs allow the focal firm to easily integrate the services of another company into their own product. As such, using an API requires the ability to modularize part of the focal firm’s product. For example, if an entrepreneur wants to provide a map feature on her product, rather than building out an entire map system and interface herself, she could simply call the Google Maps API to provide the functionality she needs in her own product. For this to work, the map component of the product must be a modular component.

The extant literature generally views modularity as the opposite of complexity (Baldwin & Clark, 2000). We assert that, holding everything else fixed, using more modular components means that the firm is not building as much proprietary complex code internally since the APIs substitute for internal code development (e.g., in the case where a map needs to do novel things and integrate novel interdependencies, using the Google API is less likely to be an option). As such, we expect firms with more novel recombination to use fewer APIs.

As outlined in Appendix Table B.10, we find that Ln API Count is generally positively correlated with the number of months the product has existed (firms add more APIs over time), but the interaction term between *Months* and *Novel Recombination* is negative ($p = 0.070$). Firms pursuing any novel recombination add one API fewer per month ($e^{(0.0005)} = -1$). We assert that this provides directional support for the increasing need to manage complexity for firms pursuing novel recombination; more novel recombination leads to increasing interdependencies that cannot be broken into modular components, such as APIs.

Table B.10: Novel Recombination and Modularity. Ordinary least squares (OLS) estimation. Robust standard errors are clustered at the firm level. p -values are shown in brackets.

	Ln API Count	Ln API Count
Months	0.021 [<0.001]	0.021 [<0.001]
Novel Recombination Binary \times Months		-0.0005 [0.070]
Observations	760 678	760 678
R^2	0.794	0.794
FE: Firm	X	X
FE: Founding Year	X	X



The Mediocrity Trap in Entrepreneurial Learning

C.1 APPENDIX

C.1.1 PROOF FOR OBSERVATION 1

Observation 1 states that given the Bayesian updating formula:

$$\mu_1 = \frac{\mu - \mu q}{1 - \mu q}$$

where μ is the prior belief and q is the initial product quality, the actor updates her prior less for lower values of q and more for higher values of q .

We aim to show that the update is more sensitive to higher values of q . To do this, we compute the derivative of μ_1 with respect to q .

DIFFERENTIATE THE UPDATING FORMULA

The updating formula is:

$$\mu_1 = \frac{\mu - \mu q}{1 - \mu q}$$

Using the quotient rule, we have:

$$\frac{\partial \mu_1}{\partial q} = \frac{(-\mu)(1 - \mu q) - \mu(1 - q)(-\mu)}{(1 - \mu q)^2}$$

Simplifying, the derivative becomes:

$$\frac{\partial \mu_1}{\partial q} = \frac{\mu(\mu - 1)}{(1 - \mu q)^2}$$

ANALYZE THE SIGN OF THE DERIVATIVE

The denominator $(1 - \mu q)^2$ is always positive since it is a square. The numerator $\mu(\mu - 1)$ depends on μ . For $0 < \mu < 1$, we have $\mu - 1 < 0$, so the numerator is negative. Thus, for $0 < \mu < 1$:

$$\frac{\partial \mu_1}{\partial q} < 0$$

This implies that as q increases, μ_1 decreases, meaning the actor updates her prior more for higher values of q .

CONCLUSION

Since $\frac{\partial \mu_1}{\partial q} < 0$, we conclude that the actor's belief update is more sensitive to higher product quality q . Therefore, the actor updates her prior less for lower values of q and more for higher values of q .

C.1.2 PROOF FOR PROPOSITION 1

Proposition 1 states: For a non-empty part of the admissible (μ, c_s, a, q) parameter space, there exist \hat{q} with $0 < \hat{q} < 1$ such that for $q < \hat{q}$, $E(\text{Sell}, \text{Sell}) > E(\text{Sell}, \text{Pivot/Quit})$ and for $q > \hat{q}$, $E(\text{Sell}, \text{Pivot/Quit}) > E(\text{Sell}, \text{Sell})$.

To prove Proposition 1, I will show that there is a $\hat{\mu}$, which is a function of c_s and a and which falls in $(0, 1)$ for at least some admissible c_s and a , and a $\tilde{q} \in [0, 1)$ such that the statement of the proposition holds for $\mu > \hat{\mu}$ and for $q > \tilde{q}$.

To find the range of values of q where $E(\text{Sell}, \text{Sell}) > E(\text{Sell}, \text{Pivot/Quit})$, it suffices to show that $E(\text{Sell}, \text{Sell}) - E(\text{Sell}, \text{Pivot/Quit})$ is strictly positive between some thresholds q_{lower} and q_{upper} , where $0 < q_{\text{lower}} < q_{\text{upper}} < 1$. Indeed, when $E(\text{Sell}, \text{Sell}) - E(\text{Sell}, \text{Pivot/Quit})$ is strictly positive, the entrepreneur will prefer to continue selling after failure rather than pivot/quit.

After some simplification, we express the difference $E(\text{Sell}, \text{Sell}) - E(\text{Sell}, \text{Pivot/Quit})$ as:

$$f(q) = (1 - \mu q) \left[\frac{\mu - \mu q}{1 - \mu q} q - c_s - a \right].$$

Since $1 - \mu q > 0$ given $\mu \in (0, 1)$ and $q \in (0, 1)$, it further follows that $f(q) > 0$ if and only if:

$$f(q) \times (1 - \mu q) > 0.$$

This allows us to simplify and solve for the roots of some $g(q)$:

$$g(q) = \frac{\mu - \mu q}{1 - \mu q} q - c_s - a$$

We rearrange it to the following quadratic expression:

$$g(q) = -\mu q^2 + (\mu + (c_s + a)\mu)q - (c_s + a).$$

This quadratic equation represents a ****downward-opening parabola**** because the coefficient of q^2 is negative ($-\mu$). This means that in the region between the two roots, $E(\text{Sell, Sell}) - E(\text{Sell, Pivot/Quit})$ is positive and the entrepreneur E will persist in selling rather than pivot or quit. Thus, to finish proving Proposition 1, it suffices to show that the quadratic has two real and distinct roots between 0 and 1.

ROOTS OF THE QUADRATIC EQUATION

Using the quadratic formula, we solve for q in the equation $g(q) = 0$. The quadratic formula is given by:

$$q = \frac{-B \pm \sqrt{B^2 - 4AC}}{2A},$$

where:

$$A = -\mu,$$

$$B = \mu(1 + c_s + a),$$

$$C = -(c_s + a).$$

Substituting these values, we get:

$$q = \frac{-\mu(1 + c_s + a) \pm \sqrt{\mu^2(1 + c_s + a)^2 - 4(-\mu)(-(c_s + a))}}{2(-\mu)}.$$

Simplifying the square root and the expression:

$$q = \frac{(1 + c_s + a) \pm \sqrt{(1 + c_s + a)^2 - \frac{4(c_s + a)}{\mu}}}{2}.$$

CONDITION FOR REAL ROOTS

To ensure that the quadratic has real roots, the discriminant must be positive. The discriminant is:

$$(1 + c_s + a)^2 - \frac{4(c_s + a)}{\mu}.$$

For real roots, this must be positive:

$$(1 + c_s + a)^2 > \frac{4(c_s + a)}{\mu}.$$

Multiplying both sides by μ gives:

$$\mu(1 + c_s + a)^2 > 4(c_s + a).$$

Dividing by $(1 + c_s + a)^2$, we obtain the condition:

$$\hat{\mu} > \frac{4(c_s + a)}{(1 + c_s + a)^2}.$$

PROVING THAT $0 < q_{\text{LOWER}} < q_{\text{UPPER}} < 1$

Next, we show that the roots q_{lower} and q_{upper} lie within the interval $(0, 1)$.

1. **Prove that $q_{\text{lower}} > 0$:**

For $q_{\text{lower}} = \frac{(1+c_s+a) - \sqrt{(1+c_s+a)^2 - \frac{4(c_s+a)}{\mu}}}{2}$, we need to show that the numerator is positive.

Since the discriminant condition ensures that the square root is less than $(1 + c_s + a)$, it follows that $q_{\text{lower}} > 0$.

2. **Prove that $q_{\text{upper}} < 1$:**

For $q_{\text{upper}} = \frac{(1+c_s+a) + \sqrt{(1+c_s+a)^2 - \frac{4(c_s+a)}{\mu}}}{2}$, we need $(1 + c_s + a) + \sqrt{(1 + c_s + a)^2 - \frac{4(c_s+a)}{\mu}} <$

2. Since the discriminant ensures that the square root term is bounded, $q_{\text{upper}} < 1$.

3. **Prove that $q_{\text{lower}} < q_{\text{upper}}$:**

The quadratic formula gives two distinct roots, with q_{lower} being the smaller root and q_{upper} the larger. Given that $\sqrt{(1 + c_s + a)^2 - \frac{4(c_s+a)}{\mu}}$ is positive, it follows that $q_{\text{lower}} < q_{\text{upper}}$.

LOWER BOUND ON Q

Without a lower bound on q , however, we cannot say that there exists a \hat{q} such that $E(\text{Sell}, \text{Sell}) > E(\text{Sell}, \text{Pivot/Quit})$ for $q < \hat{q}$ because once $q < q_{\text{lower}} < \hat{q}$, we again have that $E(\text{Sell}, \text{Pivot/Quit}) > E(\text{Sell}, \text{Sell})$, so any such q would violate the result of the proposition. As such, order for Proposition 1 to hold, we restrict to the part of the parameter space with $q > q_{\text{lower}} = \tilde{q}$

CONCLUSION

Thus, we have proven that there exist two thresholds q_{lower} and q_{upper} such that:

- For $q \in (q_{\text{lower}}, q_{\text{upper}})$, $E(\text{Sell}, \text{Sell}) > E(\text{Sell}, \text{Pivot/Quit})$, meaning the entrepreneur prefers to persist in selling.

- For $q < q_{\text{lower}}$ or $q > q_{\text{upper}}$, $E(\text{Sell, Sell}) \leq E(\text{Sell, Pivot/Quit})$, meaning the entrepreneur prefers to quit or pivot.

The condition $\hat{\mu} > \frac{4(c_s+a)}{(1+c_s+a)^2}$ ensures that q_{lower} and q_{upper} are real and lie within the interval $(0, 1)$, completing the proof of **Proposition 1**.

C.1.3 COMPARATIVE STATICS: RANGE OF INITIAL PRODUCT QUALITY SUBJECT TO THE MEDIOCRITY TRAP

In this section, I take the range of values of initial product quality q where the entrepreneur E will be subject to the mediocrity trap and then determine how the size of this region is impacted by m , c_s , and a . To start, and repeating some of what was done in the proof previously, the region can be defined by taking the difference between these two roots, ΔQ , is:

$$\begin{aligned} \Delta Q &= q_{\text{upper}} - q_{\text{lower}} \\ &= \frac{\left(1 + a + c_s + \sqrt{(1 + a + c_s)^2 - \frac{4(a+c_s)}{m}}\right) - \left(1 + a + c_s - \sqrt{(1 + a + c_s)^2 - \frac{4(a+c_s)}{m}}\right)}{2} \end{aligned}$$

Simplifying:

$$\Delta Q = \sqrt{(1 + a + c_s)^2 - \frac{4(a + c_s)}{m}}$$

The discriminant D in the expression for ΔQ is given by:

$$D = (1 + a + c_s)^2 - \frac{4(a + c_s)}{m}$$

For ΔQ to be a real number, D must be non-negative:

$$(1 + a + c_s)^2 - \frac{4(a + c_s)}{m} \geq 0$$

Rearranging to isolate m :

$$m \geq \frac{4(a + c_s)}{(1 + a + c_s)^2}$$

If this condition is not met, this means that there are no values of initial product quality q where E will fall into the mediocrity trap.

Given that D must be positive for ΔQ to be defined (or in other words for the mediocrity trap to be possible), I can simplify the analysis of comparative statics by taking partial derivatives of D (ignore the square root) but make it subject to the constraint that D must be positive. I now consider the partial derivatives of D with respect to m , c_s , and a , and interpret their implications.

The partial derivative with respect to m is:

$$\frac{\partial D}{\partial m} = \frac{4(a + c_s)}{m^2}$$

Given that D must be non-negative, this derivative indicates that as m increases, D increases. In particular, looking back at ΔQ , we see that as m increases, the term $\frac{4(a+c_s)}{m}$ decreases, making it easier for the discriminant to remain non-negative. Therefore, ΔQ is more likely to be defined, and the region where ΔQ is defined expands. In other words, a larger greater initial certainty a market exists corresponds to a larger range of values of initial product quality where the entrepreneur E will fall into the mediocrity trap and persist in spite of negative signals.

The partial derivative of D with respect to c_s is:

$$\frac{\partial D}{\partial c_s} = 2(1 + a + c_s) - \frac{4}{m}$$

Given that D must be non-negative, we can assume that $\frac{4}{m} < 2(1 + a + c_s)$ and just examine the impact of changing values of c on $2(1 + a + c_s)$. This suggests that as c increases, D decreases. Again, looking back at ΔQ , we see that as c_s increases, the numerator $4(a + c_s)$ increases faster than the denominator $(1 + a + c_s)^2$, which increases the threshold for m needed to keep the discriminant non-negative. This makes it more difficult for ΔQ to be defined, causing the region where ΔQ is defined to shrink. In other words, a higher cost of trying to sell corresponds to a smaller range of values of initial product quality where the entrepreneur E will fall into the mediocrity trap and persist in spite of negative signals.

The partial derivative of D with respect to a is:

$$\frac{\partial D}{\partial a} = 2(1 + a + c_s) - \frac{4}{m}$$

This derivative is structurally identical to the derivative with respect to c_s . It suggests that as a increases, D decreases. In other words, a more attractive alternative option to which to pivot corresponds to a smaller range of values of initial product quality where the entrepreneur E will fall into the mediocrity trap and persist in spite of negative signals.

In summary, the range of values of initial product quality q for which E will fall into the mediocrity trap increases with m and decreases with c_s and a ; E is more likely to fall into the mediocrity trap when she has a higher prior belief that the market exists, when cost of selling is low, and when the expected value of her alternative option from pivoting is low.

C.1.4 MODEL LIMITATIONS AND EXTENSIONS

The model presented in this paper deliberately simplifies entrepreneurial decision-making to highlight the mediocrity trap mechanism. While these simplifications help maintain analytical tractability and clarify the core insight, they also limit the model. Several extensions could enrich the analysis and

provide additional insights.

MULTI-PERIOD DYNAMICS: The model indirectly captures product learning through the Δq parameter and the choice between immediate market entry versus development first. However, a multi-period extension would explicitly model how early market feedback could increase the efficiency of subsequent development efforts, potentially making Δq a function of prior market experience rather than a fixed parameter.

COMPETITION: The model implicitly incorporates competition through several parameters: μ reflects knowledge of the competitive landscape in market existence beliefs, while q incorporates the quality threshold needed to win against alternatives. An explicit competition extension would model how these parameters dynamically respond to competitor actions rather than treating them as fixed values. **Network effects:** The model could incorporate how each successful sale automatically increases product quality, creating self-reinforcing adoption dynamics.

NETWORK EFFECTS: Many products, particularly in technology and platform markets, exhibit network effects where the value of the product increases as more customers use it. The current model treats product quality as either fixed or improved only through deliberate development investment, without considering how successful sales themselves might increase product quality.

Incorporating network effects would model how each successful sale automatically increases product quality for subsequent customers, creating self-reinforcing adoption dynamics. This extension would be particularly valuable for analyzing products like the APIs studied in the empirical section, where more users often means greater value for all users. It would show how entrepreneurs facing strong network effects might rationally prioritize early market entry with lower-quality products, as the sales process itself generates quality improvements that make future sales more likely. Network

effects could potentially mitigate the mediocrity trap by creating a path where initial quality, while not great, is sufficient to start a positive feedback loop.

MARKET EVOLUTION: The model assumes market conditions remain static, with market existence treated as a fixed but unknown state. All changes to purchase probability stem from endogenous product quality improvements rather than exogenous market changes.

In reality, markets evolve through complementary innovation, shifting customer needs, and broader ecosystem development. A time-varying market parameters extension would capture how markets emerge and evolve independently of the entrepreneur's actions. This would allow modeling of scenarios where initially non-existent markets materialize over time or where early markets disappear due to external factors. Such an extension would be particularly valuable for analyzing timing strategies in emerging technology markets, where being too early ("ahead of the market") or too late can both lead to failure. It would help identify conditions under which patience in market development might be more valuable than rapid pivoting.

SEGMENT HETEROGENEITY: The current model treats the market as homogeneous, with a single quality threshold determining sales probability and each potential customer viewed as identical. This simplification doesn't account for how different customer segments might have varying quality requirements, willingness to pay, or specific needs.

Incorporating heterogeneous market segments would model how entrepreneurs might strategically target different types of customers as part of their learning strategy. This extension would show how entrepreneurs could sequence market entry across segments, perhaps targeting more accessible early adopters before addressing more demanding mainstream customers. It would also demonstrate how segment-specific feedback might inform product development priorities differently. Such an extension would connect to existing literature on crossing the chasm between early adopters and main-

stream markets, showing how the mediocrity trap might manifest differently when entrepreneurs can choose which market segments to target first.

QUALITY UNCERTAINTY: The model assumes entrepreneurs have perfect information about their product's quality. Quality is treated as an objective parameter that is common knowledge to both the entrepreneur and potential customers.

In reality, entrepreneurs often face uncertainty about their own product's quality and how it will be perceived by customers. They may over or underestimate their product's capabilities relative to customer needs. Incorporating quality uncertainty would model how entrepreneurs form and update beliefs about their product quality through market feedback. This extension would show how quality misestimation affects persistence decisions - overconfidence might exacerbate the mediocrity trap by causing entrepreneurs to attribute failures entirely to market factors rather than product inadequacies, while underconfidence might lead to premature abandonment of viable opportunities.

C.1.5 ALTERNATIVE MEASURES OF PERSISTENCE IN SELLING AGAIN

While the main analysis uses the number of new salespeople in the subsequent year to proxy a firm's persistence in selling again, I also explore two additional alternative measures of selling again. I find consistent results with both of these alternative measures, adding validity and additional empirical support for Hypothesis 6.

SELL AGAIN: TOTAL SALESPEOPLE MONTHS *Sell Again: Total Salespeople* is a leading variable equal to the total number of months salespeople were employed by the firm in the following year.

SELL AGAIN: SALESPEOPLE ATTRITION RATIO *Sell Again: Churned Salespeople* is a leading variable equal to the total number of salespeople employed by the firm in the current year who leave the

firm in the following year normalized relative to the total number of salespeople in a standard calculation of churn.

Additional support for Hypothesis 6 is provided by the results for *Total Salespeople Months* and *Salespeople Attrition* in Table C.1. The coefficient for *CAC* on *Total Salespeople Months* is positive and significant ($\beta = 0.242, p < 0.01$), indicating that a 100% increase in *CAC* is associated with an 18.3% increase in total salespeople months. This suggests that firms not only hire new salespeople but also retain existing ones when faced with high *CAC*, further demonstrating persistence in selling efforts.

Table C.1: Alternative Measures of Persistence in Selling Again. Ordinary least squares (OLS) estimation. Standard errors (in parentheses) are clustered at the year level. *p*-values are shown in brackets.

	Ln Total Salespeople Months	Ln Salespeople Attrition
Ln Product Development	0.080 (0.020) [0.005]	-0.031 (0.020) [0.170]
Ln CAC	0.242 (0.025) [0.000]	-0.345 (0.025) [0.000]
Ln Product Development × Ln CAC	-0.035 (0.005) [0.000]	0.039 (0.005) [0.000]
Ln Funding	0.117 (0.035) [0.013]	-0.082 (0.029) [0.026]
Ln Total Employees	0.377 (0.032) [0.000]	-0.385 (0.045) [0.000]
(Intercept)		
Observations	9343	9343
R^2	0.800	0.797
FE: Year	X	X
FE: Firm	X	X

Moreover, the negative and significant coefficient for *CAC* on *Salespeople Attrition* ($\beta = -0.345, p < 0.01$) indicates a substantial decrease in attrition by about 21.3% with a doubling of *CAC*. This suggests that firms are more likely to retain their salespeople despite higher acquisition costs, providing additional evidence of persistence in selling efforts when faced with negative market signals.

In summary, these results strongly support Hypothesis 1, demonstrating that enterprise software startups tend to persist in their selling efforts by increasing total salespeople months and reducing

salespeople attrition when faced with high customer acquisition costs, particularly when they have lower product development investment.

C.1.6 SELECTION CONCERNS

This study relies on the integration of multiple data sources, each with potential selection concerns that warrant careful consideration. I address these concerns and their implications for the interpretation of the results below.

PITCHBOOK Pitchbook is widely regarded as the most comprehensive startup database available at scale. However, it may omit firms that never publicize their founding or launch any type of product. As a result, the sample likely overrepresents more established or public-facing startups. This potential selection bias suggests that the results should be interpreted with the boundary condition that they apply to startups that have achieved a certain level of visibility or development.

LINKEDIN LinkedIn profiles serve as the source of employee information. While LinkedIn has high penetration in the professional world, especially in the tech sector, selection concerns could arise if there are systematic differences in the small group of startup employees without profiles. As such, the results may be most applicable to firms with employees who are active on professional networking platforms.

BUILTWITH As noted by [Stroube & Dushnitsky \(2023\)](#), BuiltWith data is most complete and robust after 2012. While the vast majority of the panel is post-2012, there are some observations starting earlier. To ensure this is not introducing any selection concerns that might bias the results, I rerun regressions with the subset of observations post-2012 and find that the results hold.

Appendix Table C.2 contains the regression analysis for *Sell Again*, subset to only include data from years after 2012, and shows that the results hold.

Table C.2: BuiltWith Selection: Post-2012. Ordinary least squares (OLS) estimation. Standard errors (in parentheses) are clustered at the year level. *p*-values are shown in brackets.

	Sell Again
Ln Product Development	0.027 (0.012) [0.070]
Ln CAC	0.016 (0.011) [0.185]
Ln Product Development × Ln CAC	−0.005 (0.003) [0.092]
Ln Funding	0.037 (0.026) [0.193]
Ln Total Employees	0.081 (0.020) [0.005]
Observations	4758
R^2	0.701
FE: Year	X
FE: Firm	X

C.1.7 COMPETITION

Competition in the enterprise software market has significant implications for both value creation and value capture. From a theoretical perspective, the presence of competitors can be viewed through two distinct lenses: value creation and value capture (Brandenburger & Stuart Jr., 1996).

VALUE CREATION: The existence of competitors in a market can serve as a signal that there is indeed a valuable market opportunity. This aligns with the concept of vicarious learning, where firms can learn about market potential by observing the actions of their competitors (Huber, 1991). Furthermore, competition can drive innovation and market expansion, potentially increasing the overall

value created in the industry (Aghion et al., 2005).

VALUE CAPTURE: While competition may signal market potential, it also presents challenges for individual firms in capturing that value. As Porter (1980) articulates in his seminal work on competitive strategy, the intensity of rivalry among existing competitors is a key force that shapes industry profitability. In markets with intense competition, firms may struggle to differentiate their offerings or maintain pricing power, potentially leading to higher customer acquisition costs (CAC) as they compete for the same pool of customers.

These two aspects of competition—its role in signaling market potential and its impact on value capture—have important implications for how we interpret the relationship between competition, product development, and CAC in this study. The presence of competitors may simultaneously increase a firm’s confidence in the market opportunity (potentially encouraging persistence) while also making it more difficult to acquire customers efficiently (potentially leading to higher CAC).

I argue that competition is inherently embedded in the market feedback captured by Customer Acquisition Cost (CAC), both in the theoretical model and in the empirical analysis, and therefore not directly measuring it is not a cause for concern because it is already accounted for indirectly. A high CAC serves as a comprehensive negative signal about product-market fit, regardless of its specific underlying causes. Whether potential customers are not purchasing due to the presence of superior competing products, the lack of a compelling value proposition, or any other market factors, the end result manifests in an elevated CAC. This metric effectively encapsulates various market forces, including competition, without necessitating their individual measurement. In the context of this study, the precise reason for a high CAC is less critical than its implications for the firm’s strategic decisions. The theory posits that firms should interpret a high CAC as a negative signal about product-market fit, irrespective of whether this is driven by direct competition or other factors affecting customer demand.

Moreover, the decision to persist in selling efforts or pivot in response to a high CAC should not

be contingent on isolating competition as a factor. If a firm faces high customer acquisition costs, it suggests that the current product offering is not sufficiently compelling to customers relative to their alternatives, which inherently includes competing products. In this scenario, the rational response would be to reassess the product-market fit rather than to intensify sales efforts. The mediocrity trap identified in this study occurs when firms with less-developed products fail to interpret this signal correctly, persisting in their sales efforts despite the negative feedback from the market. This persistence is problematic regardless of whether the high CAC is primarily driven by strong competition or other factors affecting product-market fit. Therefore, while competition is undoubtedly an important market force, its direct measurement is not necessary for the core arguments and findings of this study. The CAC metric serves as a sufficient proxy for the overall market response, including competitive effects, and allows for a focused examination of how firms interpret and act on this aggregate signal in relation to their product development investments.

EMPIRICAL ANALYSIS OF COMPETITION

In addition to outlining the above theoretical reasons why competition is captured in CAC, I also construct a measure of competition and test its impact directly.

For each firm in the sample who lists a set of competitors in Pitchbook, I use this data to create a directed graph with each firm pointing to its identified competitors. From here, there are many different approaches to break this directed graph into clusters. These different approaches balance considerations such as whether to allow a company to appear in multiple clusters, maximum cluster size, maximum computing time, memory requirements, etc. Given the various trade-offs across these different dimensions, I employ two different techniques to create clusters. These techniques include:

LEIDEN ALGORITHM (Traag et al., 2019) The Leiden Algorithm is often used for single-cell RNA sequencing data analysis. It operates by optimizing a quality function that measures the significance

of node assignments to clusters. Leiden starts with an initial clustering and iteratively evaluates the quality of the clusters by considering the benefits of moving nodes between clusters. If moving nodes enhances cluster quality, the algorithm makes that adjustment. This iterative process continues until no further improvements can be made.

INFOMAP ALGORITHM (Rosvall & Bergstrom, 2008) The Infomap Algorithm analyzes the network's structure and seeks to minimize the information needed to describe it. It starts by exploring the network using random walks, tracking how information flows between nodes. The algorithm identifies clusters by finding regions where information can be efficiently compressed. It optimizes these clusters by iteratively merging or splitting them to minimize the description length.

The Leiland algorithm is scalable, making it suitable for analyzing large networks. I adjusted the resolution parameter to 28, which keeps the maximum cluster size at a reasonable number. The Infomap Algorithm is also computationally efficient: it is fast and generates reasonable cluster sizes.

After defining each cluster of competitors using the approaches outlined above, I subset to the firms who are competitors within the sample. I then use the founding year of each firm within the cluster to define a time-variant measure of the number of competitors within each cluster at any given point in time. This allows for a dynamic proxy of the level of competition. *Leiden Competitors* is equal to the total number of competitors the focal firm has in a given year using the Leiden method. *Infomap Competitors* is equal to the total number of competitors the focal firm has in a given year using the Infomap method.

Appendix Table C.3 presents regression results incorporating these different measures of competition for the *Sell Again* regression specification. For both competition measures, the baseline term for competition is positive and significant ($\beta = 0.492, p = 0.010$ for Leiden; $\beta = 0.436, p = 0.004$ for Infomap). This suggests that increased competition in the current year is associated with hiring more new salespeople in the subsequent year; firms may respond to competitive pressure by intensifying

sales efforts.

The interaction term between competition and product development is negative, though not statistically significant at conventional levels ($\beta = -0.018$, $p = 0.170$ for Leiden; $\beta = -0.019$, $p = 0.132$ for Infomap). This trend suggests that firms with more developed products may be less likely to respond to increased competition by adding more new salespeople, consistent with the idea that higher-quality products may better withstand competitive pressures.

Importantly, the interaction term between CAC and product development remains negative and significant when controlling for competition ($\beta = -0.015$, $p = 0.076$ for Leiden; $\beta = -0.017$, $p = 0.073$ for Infomap). This robustness supports the main findings regarding the mediocrity trap, indicating that the relationship between product development, CAC, and persistence in selling efforts holds even when accounting for competitive dynamics.

I also run an alternative regression specification where I include fixed effects for the time-varying number of competitors and similarly find that the effect holds. Table C.4 shows this for the *Sell Again* regression specification. Note that the interaction term is still negative and significant, suggesting the evidence for persistence in the mediocrity trap is robust to competition.

C.1.8 MULTIPLE PRODUCTS

The potential impact of multiple products on the analysis can be considered in two scenarios: sequential products and simultaneous products.

SEQUENTIAL PRODUCTS

Multiple products can occur sequentially if a firm introduces a new product and stops selling its original product. In this case, it is as if the firm is simply starting over. The prior product should not impact the CAC for the new product or the firm's propensity to hire new salespeople next year to

Table C.3: Persistence in Selling Again and Competition Variables. Ordinary least squares (OLS) estimation. Standard errors (in parentheses) are clustered at the year level. *p*-values are shown in brackets.

	Sell Again	Sell Again
Ln Product Development	0.068 (0.029) [0.051]	0.070 (0.029) [0.045]
Ln CAC	0.023 (0.037) [0.554]	0.027 (0.036) [0.477]
Ln Leiden Competitors	0.492 (0.139) [0.010]	
Ln Funding	0.080 (0.029) [0.029]	0.076 (0.029) [0.036]
Ln Total Employees	0.175 (0.065) [0.031]	0.185 (0.066) [0.027]
Ln Product Development × Ln CAC	−0.015 (0.007) [0.076]	−0.017 (0.008) [0.073]
Ln Product Development × Ln Leiden Competitors	−0.018 (0.012) [0.170]	
Ln CAC × Ln Leiden Competitors	−0.003 (0.012) [0.779]	
Ln Product Development × Ln CAC × Ln Leiden Competitors	0.003 (0.003) [0.392]	
Ln Infomap Competitors		0.436 (0.103) [0.004]
Ln Product Development × Ln Infomap Competitors		−0.019 (0.011) [0.132]
Ln CAC × Ln Infomap Competitors		−0.005 (0.012) [0.698]
Ln Product Development × Ln CAC × Ln Infomap Competitors		0.003 (0.003) [0.333]
Observations	2710	2710
<i>R</i> ²	0.752	0.751
FE: Year	X	X
FE: Firm	X	X

Table C.4: Persistence in Selling Again and Competition Fixed Effects. Ordinary least squares (OLS) estimation. Standard errors (in parentheses) are clustered at the year level. p -values are shown in brackets.

	Sell Again	Sell Again
Ln Product Development	0.024 (0.013) [0.118]	0.028 (0.015) [0.101]
Ln CAC	0.014 (0.023) [0.556]	0.019 (0.022) [0.422]
Ln Product Development \times Ln CAC	-0.009 (0.004) [0.068]	-0.011 (0.004) [0.020]
Ln Funding	0.076 (0.028) [0.029]	0.074 (0.029) [0.036]
Ln Total Employees	0.173 (0.062) [0.027]	0.172 (0.062) [0.028]
Observations	2710	2710
R^2	0.761	0.762
FE: Year	X	X
FE: Firm	X	X
FE: Leiden Competitors	X	
FE: Infomap Competitors		X

try to sell the new product. As such, there should not be a systematic bias. The main distinction is that the cumulative product investment value would be inflated; the firm would have a larger value for *Product Development* even though the new product it was working on and trying to sell was much less developed. This would bias the results against my main finding. I find that firms with lower levels of product investment are more likely to persist and hire new salespeople. However, in the case of multiple sequential products, firms with lower levels of product investment would actually appear to have high levels of product investment. Given that the noise introduced by multiple sequential products would work against the result, accounting for sequential products should only strengthen the results.

However, in practice, it is very uncommon for a firm to completely abandon a prior product such that it would even be fully starting over on product development. For example, Slack started as game called Glitch with an internal communication tool that eventually became Slack's flagship product after the game shutdown. Or Twitter began as a side project within a podcasting platform called Odeo. When Apple launched iTunes with podcast support, Odeo's core business became obsolete so the team expanded the side project they had been working on—a microblogging service that eventually became Twitter. These examples illustrate how even subsequent new products do benefit from the prior development and so the product development number is likely not as overinflated as might initially seem could be the case.

SIMULTANEOUS PRODUCTS

Multiple products can also occur simultaneously if a firm is developing and trying to sell two or more products at the same time. When firms develop and sell multiple products concurrently, it introduces measurement challenges for some of the key variables. The presence of multiple products likely adds noise to the measures of product development investment and Customer Acquisition Cost (CAC). For instance, the product development measure might overestimate the investment in any single product, as it captures total investment across all products. Similarly, CAC might be misattributed be-

tween products when sales efforts or customer acquisitions aren't clearly delineated; CAC will reflect the average across all a firm's products.

Importantly, this measurement error is likely to be non-systematic. There is no clear reason to believe that firms with multiple products would consistently over or under invest in product development relative to single-product firms, or that they would systematically have higher or lower CAC; the different products could be easier or harder to sell. As such, the presence of multiple products likely introduces noise rather than bias into the analysis.

This measurement error typically leads to attenuation bias, making the estimates more conservative. In other words, it becomes more difficult to detect the true relationship between product development, CAC, and persistence in selling efforts. The fact that there are still significant relationships in line with the hypotheses, despite this added noise, strengthens confidence in the results.

EMPIRICAL IMPACT OF MULTIPLE PRODUCTS

The empirical analysis in Appendix Table C.5 addresses the potential impact of multiple products on the main findings. The regression includes a time-varying indicator variable, *Single Product*, which equals 1 if the firm has a single product and 0 otherwise. The results show a significant negative main effect for *Single Product* ($\beta = -0.333, p = 0.020$), suggesting that firms with a single product tend to hire fewer new salespeople compared to those with multiple products; firms with multiple products might require larger sales forces.

The interaction between *Single Product* and CAC is negative and significant ($\beta = -0.088, p = 0.046$), indicating that single-product firms are less likely to respond to high CAC by hiring additional salespeople compared to multi-product firms.

Crucially, the interaction term between CAC and product development remains negative and significant when controlling for single versus multiple products ($\beta = -0.025, p = 0.006$). Furthermore, three-way interaction between Product Development, CAC, and *Single Product* is positive and

Table C.5: Single vs Multiple Products Over Time. Ordinary least squares (OLS) estimation. Standard errors (in parentheses) are clustered at the year level. *p*-values are shown in brackets.

	Sell Again
Ln Product Development	0.048 (0.034) [0.202]
Ln CAC	0.100 (0.039) [0.038]
Single Product	-0.333 (0.111) [0.020]
Ln Funding	0.099 (0.027) [0.009]
Ln Total Employees	0.109 (0.018) [<0.001]
Ln Product Development × Ln CAC	-0.025 (0.006) [0.006]
Ln Product Development × Single Product	-0.001 (0.028) [0.977]
Ln CAC × Single Product	-0.088 (0.036) [0.046]
Ln Product Development × Ln CAC × Single Product	0.021 (0.007) [0.024]
Observations	8633
R^2	0.629
FE: Year	X
FE: Firm	X

significant ($\beta = 0.021, p = 0.024$). This suggests that the mediocrity trap effect (i.e., the negative interaction between Product Development and CAC) is less pronounced for single-product firms and more pronounced for multi-product firms.

C.1.9 OTHER CHANNELS FOR CUSTOMER ACQUISITION

There could be a concern that the CAC measure, based on the ratio of salespeople to new customers, may not capture customer acquisition efforts through other channels such as advertising. This could potentially lead to underestimation of true customer acquisition costs for firms heavily reliant on advertising. However, I argue that this limitation does not significantly undermine the findings and may, in fact, make the results more conservative.

First, if firms with less-developed products resort to heavy advertising due to challenges with direct sales, this could result in these firms appearing to have low CAC in our measure. This would work against the hypothesis, making it more difficult to detect the relationship between low product development and persistence in the face of high CAC. The fact that there is still a significant result suggests the findings are robust.

Second, in the B2B SaaS context of our study, direct sales often play a crucial role, potentially limiting the prevalence of firms relying solely on advertising for customer acquisition. Moreover, firms successful enough to invest heavily in advertising are likely to also invest in salespeople to capitalize on generated leads, which our measure would capture.

Finally, even if a firm acquires many customers with few salespeople (resulting in low measured CAC), this still indicates a degree of product-market fit, as customers are choosing to buy with minimal direct sales effort. This aligns with the overall theory about the importance of product development in achieving market fit.

EMPIRICAL ANALYSIS OF ADVERTISING

To address the potential impact of advertising on customer acquisition, I first examine the relationship between traditional sales efforts and online advertising investments. Appendix Figure C.1 illustrates the correlation between the number of salespeople and the number of paid keywords used by firms in our sample.




Figure C.1: Binned Scatter of Salespeople and Paid Keywords. This scatter plot shows the relationship between the natural log of the number of salespeople (x -axis) and the natural log of the number of paid keywords (y -axis) for firms in our sample. The positive trend suggests that firms investing more in traditional sales forces also tend to invest more in online advertising through paid keywords.

The positive correlation observed in Figure C.1 suggests that firms often pursue both traditional sales and online advertising strategies in tandem, rather than treating them as substitutes. This complementary approach to customer acquisition motivates our investigation into how advertising efforts

might influence the relationship between product development, Customer Acquisition Cost (CAC), and persistence in selling efforts.

To further explore this relationship, I introduce a measure of advertising effort based on paid keywords. *Paid Keywords* is equal to the total keywords a website is buying in Google Ads for ads that appear in paid search results. This measure serves as a proxy for a firm's investment in online advertising, which could be an alternative or complementary channel to direct sales efforts.

To incorporate this advertising measure into our analysis, I construct two alternative CAC metrics:

1. *CAC Paid Keywords* is equal to the total number of paid keywords in a given year divided by the total new customers acquired in that year.
2. *CAC Combined* is equal to the total number of salespeople plus the total number of paid keywords in a given year, all divided by the total new customers acquired by a given firm in a given year.

Appendix Table C.6 presents the results of regression analyses incorporating these new measures. The first column adds a control for Paid Keywords while the second and third columns introduce the alternative CAC calculations.

The results show that the number of Paid Keywords is positively associated with hiring new salespeople in the subsequent year ($\beta = 0.026, p = 0.007$). This suggests that firms investing more in online advertising also tend to expand their sales force, potentially indicating a complementary relationship between these two customer acquisition strategies. Furthermore, the interaction between Product Development and CAC remains negative and significant.

When examining the impact of CAC Paid Keywords, we observe a positive and significant relationship with new salesperson hiring ($\beta = 0.021, p = 0.031$). The interaction term between Product Development and CAC Paid Keywords is negative and marginally significant ($\beta = -0.003$,

Table C.6: Paid Keywords. Ordinary least squares (OLS) estimation. Standard errors (in parentheses) are clustered at the year level. p -values are shown in brackets.

	Sell Again	Sell Again	Sell Again
Ln Product Development	0.031 (0.011) [0.023]	0.027 (0.011) [0.042]	0.034 (0.011) [0.018]
Ln CAC	0.018 (0.010) [0.112]		
Ln Funding	0.043 (0.025) [0.132]	0.047 (0.026) [0.111]	0.049 (0.027) [0.106]
Ln Total Employees	0.085 (0.018) [0.002]	0.095 (0.023) [0.004]	0.091 (0.019) [0.002]
Ln Paid Keywords	0.026 (0.007) [0.007]		
Ln Product Development \times Ln CAC	-0.006 (0.002) [0.032]		
Ln CAC Paid Keywords		0.021 (0.008) [0.031]	
Ln Product Development \times Ln CAC Paid Keywords		-0.003 (0.001) [0.052]	
Ln CAC Combined			0.016 (0.009) [0.115]
Ln Product Development \times Ln CAC Combined			-0.004 (0.002) [0.096]
Observations	5495	5495	5495
R^2	0.694	0.694	0.693
FE: Year	X	X	X
FE: Firm	X	X	X

$p = 0.052$). This pattern is consistent with our main findings, suggesting that the mediocrity trap phenomenon persists even when considering advertising efforts.

The analysis using CAC Combined yields similar results. The main effect of CAC Combined on new salesperson hiring is positive but not statistically significant ($\beta = 0.016, p = 0.115$). However, the interaction between Product Development and CAC Combined remains negative and marginally significant ($\beta = -0.004, p = 0.096$). This further supports the robustness of our main findings.

Importantly, across all specifications, the interaction between Product Development and the various CAC measures remains negative and at least marginally significant. This consistency provides strong support for the mediocrity trap hypothesis, even when accounting for advertising efforts.

These results suggest that while advertising does play a role in customer acquisition, it does not fundamentally alter the relationship between product development, CAC, and persistence in selling efforts. The mediocrity trap phenomenon appears to hold even when considering alternative channels for customer acquisition, reinforcing the importance of product quality in interpreting and responding to market signals.

C.1.10 STRATEGY SHIFT

Changes in a firm's strategy over time could potentially influence some variables of interest, including product development investment, customer acquisition costs, and persistence in sales efforts. To account for this, I use changes in CEO as a proxy for significant shifts in a firm's strategy over time, based on the assumption that each CEO has a unique strategic approach and that boards often replace a firm's CEO when implementing a major strategic change.

Appendix Table C.7 presents empirical results addressing potential strategy shifts, using CEO changes as a proxy. The analyses include a *New CEO* indicator (equal to 1 if there is a new CEO in the current year and 0 otherwise) and CEO fixed effects to account for time-invariant characteristics of each CEO's strategy. The main effects and interactions involving the *New CEO* variable are

not statistically significant at conventional levels. However, the interaction between *Product Development* and *CAC* remains negative and significant ($\beta = -0.006$, $p = 0.083$ without CEO fixed effects; $\beta = -0.006$, $p = 0.068$ with CEO fixed effects). This suggests that the main findings are robust to controlling for strategic shifts associated with CEO changes.

Table C.7: Strategic Shifts Via CEO Change. Ordinary least squares (OLS) estimation. Standard errors (in parentheses) are clustered at the year level. p -values are shown in brackets.

	Sell Again	Sell Again
Ln Product Development	0.049 (0.012) [0.004]	0.049 (0.012) [0.004]
Ln CAC	0.017 (0.011) [0.163]	0.017 (0.011) [0.160]
New CEO	-0.054 (0.077) [0.500]	
Ln Funding	0.095 (0.030) [0.015]	0.100 (0.030) [0.013]
Ln Total Employees	0.110 (0.018) [<0.001]	0.112 (0.017) [<0.001]
Ln Product Development \times Ln CAC	-0.006 (0.003) [0.083]	-0.006 (0.003) [0.068]
Ln Product Development \times New CEO	0.020 (0.016) [0.263]	
Ln CAC \times New CEO	0.034 (0.038) [0.404]	
Ln Product Development \times Ln CAC \times New CEO	-0.010 (0.006) [0.162]	
Observations	9343	9343
R^2	0.627	0.628
FE: Year	X	X
FE: Firm	X	X
FE: CEO		X

C.1.11 CUSTOMER SIZE

The size of customers to which a firm sells may impact its approach to product development, customer acquisition, and salesperson hiring. Larger customers might require different product features, sales strategies, etc. compared to smaller customers, potentially influencing the relationships observed in the main analysis.

To account for this, I use Google PageRank as a proxy for customer size. PageRank is an algorithm used by Google to rank websites in their search engine results, with higher ranks generally corresponding to more established and larger organizations. PageRank data is included in the BuiltWith data. However, it only includes PageRank up to 1,000,000 so very small customers are not ranked.

PageRank is a continuous variable equal to the average new customer's homepage Google PageRank in a given period, with 1 being the highest PageRank possible for any individual customer (which I argue is associated with a large customer size) going all the way up to a potential individual customer homepage rank of 100,000,000 (which I argue is associated with a smaller customer size).

Top 1000 Rank is a binary indicator equal to 1 if a firm's new customers' homepages in a given period average to a PageRank less than or equal to 1000 and 0 otherwise.

Appendix Table C.8 presents empirical results addressing the potential impact of customer size on the main findings. In the continuous PageRank specification, the three-way interaction between Product Development, CAC, and PageRank is not statistically significant. However, the interaction between Product Development and PageRank is positive and significant, suggesting that firms with more-developed products are more likely to acquire larger customers.

The binary *Top 1000 Rank* specification shows a significant negative main effect ($\beta = -1.308$, $p = 0.001$), indicating that firms targeting larger customers tend to hire fewer new salespeople. The interaction between Product Development and *Top 1000 Rank* is positive and significant ($\beta = 0.552$, $p = 0.004$), suggesting that firms with more developed products are more likely to target and acquire

Table C.8: PageRank Analysis: Impact of Customer Size. Ordinary least squares (OLS) estimation. Standard errors (in parentheses) are clustered at the year level. *p*-values are shown in brackets.

	Sell Again	Sell Again
Ln Product Development	0.049 (0.018) [0.029]	0.050 (0.012) [0.004]
Ln CAC	0.004 (0.024) [0.868]	0.018 (0.011) [0.138]
PageRank	0.000 (0.000) [0.843]	
Ln Funding	0.073 (0.027) [0.032]	0.095 (0.030) [0.015]
Ln Total Employees	0.193 (0.051) [0.007]	0.110 (0.018) [<0.001]
Ln Product Development × Ln CAC	-0.010 (0.005) [0.091]	-0.007 (0.003) [0.071]
Ln Product Development × PageRank	0.000 (0.000) [0.004]	
Ln CAC × PageRank	0.000 (0.000) [0.890]	
Ln Product Development × Ln CAC × PageRank	0.000 (0.000) [0.168]	
Top 1000 Rank		-1.308 (0.242) [0.001]
Ln Product Development × Top 1000 Rank		0.552 (0.133) [0.004]
Ln CAC × Top 1000 Rank		0.244 (0.171) [0.196]
Observations	3482	9343
R^2	0.772	0.627
FE: Year	X	X
FE: Firm	X	X

larger customers.

Importantly, the key interaction between *Product Development* and *CAC* remains negative and significant in both specifications ($\beta = -0.010$, $p = 0.091$ for continuous *PageRank*; $\beta = -0.007$, $p = 0.071$ for binary *Top 1000 Rank*). This indicates that the main findings regarding the mediocrity trap are robust to controlling for customer size.

C.1.12 PRICE

Given the inclusion of firm fixed effects in my analysis, price only matters if it changes over time. My primary concern is that if a company raises its price, this would also raise Customer Acquisition Cost (CAC), but the increase in CAC might be justified based on the increase in revenue associated with the price increase. This situation could potentially confound my interpretation of high CAC as a negative market signal.

However, in the context of B2B SaaS, I argue that firms are incentivized to price at customers' willingness to pay (WTP) as a baseline. Given that variable costs are essentially zero, firms are motivated to price their products to match each customer's WTP rather than lose a sale. Custom pricing agreements are the norm in this industry, which reflects this pricing strategy—firms can tailor their prices to match each customer's willingness to pay. If firms are simply pricing at the customers' WTP, then I contend that changes in pricing are not a strategic decision that could be driving CAC.

Instead, I identify competition and customer size as key factors that may drive a gap between willingness to pay and price. When competition is high, firms must price strategically rather than simply pricing at WTP. In highly competitive markets, firms might need to lower prices, potentially leading to higher CAC as more sales effort is needed to convince customers to choose their product over competitors'. This relationship between competition, pricing, and CAC further justifies my use of competition as a relevant proxy for strategic pricing decisions. As I noted in Appendix Section C.1.7, my results hold when accounting for the impact of competition, suggesting that my findings

are robust to these competitive pricing dynamics.

I also consider customer size as another factor that could drive a gap between price and WTP. Larger customers might have more bargaining power; conversely, smaller customers might have less negotiating leverage. This difference in bargaining power could affect price, as larger customers might be able to demand a more attractive price. My analysis in Appendix Section C.1.11 controls for customer size using Google PageRank as a proxy. The results remain consistent when accounting for customer size, further supporting the robustness of my findings to price concerns.

I also note that if firms with less-developed products are more likely to compete on price (potentially leading to higher CAC), this would actually strengthen my study's findings about the relationship between product development, CAC, and persistence in selling efforts. The "mediocrity trap" I identify would be even more pronounced if firms with lower-quality products are forced to price more aggressively to acquire customers.

I acknowledge that my dataset doesn't include direct price information, which limits my ability to control for price changes explicitly. However, I argue that the combination of firm fixed effects, competition controls, and customer size considerations helps mitigate concerns about price effects confounding my results. The consistency of my findings across these various specifications suggests that my core insights about the relationship between product development, CAC, and firm behavior are robust to potential pricing effects.

For future research, I suggest that incorporating more detailed pricing data could further disentangle the effects of pricing strategies from other factors affecting CAC and firm behavior. This could provide additional insights into how pricing decisions interact with product quality and market dynamics in the B2B SaaS context.

C.1.1.3 HIGH CAC AS A STRATEGIC INVESTMENT

One potential challenge to this study's interpretation of high Customer Acquisition Cost (CAC) as a negative signal is the possibility that firms strategically invest in high CAC for future benefits. For instance, in markets with strong network effects, firms might accept high initial CAC to build a user base, anticipating lower acquisition costs once network effects are established.

To empirically examine this possibility, I investigated whether firms with high CAC in the current year experience reductions in CAC in subsequent years. I defined *Leading CAC* as the customer acquisition cost in the subsequent year and analyzed its relationship with current *CAC* and *Product Development*.

This analysis reveals that CAC in the current year is positively correlated with CAC in the subsequent year ($\beta = 0.569, p < 0.001$, see Appendix Table C.9). This finding challenges the notion that high CAC consistently functions as a strategic investment for future cost reductions.

Interestingly, I find a negative interaction between current CAC and product development in predicting future CAC ($\beta = -0.020, p = 0.105$). This suggests that firms with higher product development investment may be better positioned to decrease CAC over time. This aligns with this study's main argument that product quality is crucial for effectively learning from and responding to market signals. To further explore the temporal dynamics, I extended this analysis to a two-year lead on CAC and found the result remain consistent.

These findings support this study's interpretation of high CAC as a generally negative signal rather than a strategic investment. They also reinforce the importance of product development in a firm's ability to improve its market position over time. Firms with higher product development investment appear better equipped to leverage their product quality to reduce customer acquisition costs in the future. However, I acknowledge that strategic high-CAC investments might still occur in specific cases or over shorter time frames not captured by the annual measures. Future research could explore

Table C.9: Leading Customer Acquisition Cost. Ordinary least squares (OLS) estimation. Standard errors (in parentheses) are clustered at the year level. *p*-values are shown in brackets.

	Ln Leading CAC		
Ln Product Development	-0.057 (0.009) [<0.001]	-0.013 (0.021) [0.565]	0.018 (0.018) [0.328]
Ln CAC	0.939 (0.012) [<0.001]	0.563 (0.055) [<0.001]	0.569 (0.058) [<0.001]
Ln Product Development × Ln CAC	-0.017 (0.003) [<0.001]	-0.019 (0.011) [0.128]	-0.020 (0.011) [0.105]
Ln Funding			-0.114 (0.022) [0.001]
Ln Total Employees			-0.054 (0.031) [0.122]
(Intercept)	0.035 (0.025) [0.156]		
Observations	9343	9343	9343
<i>R</i> ²	0.602	0.738	0.739
FE: Year		X	X
FE: Firm		X	X

more granular temporal patterns or identify specific market conditions where high CAC might function as an effective strategic investment.

C.1.14 LONG-TERM EFFECTS

To examine the persistence of the main findings over time, I conducted additional analyses using two- and three-year leads on our dependent variable, Sell Again. Appendix Table C.10 presents these results. For the two-year lead, the main effect of *Product Development* remains positive and significant ($\beta = 0.231, p = 0.005$), indicating that firms with higher product development investment continue to hire more salespeople two years later. However, the interaction between *Product Development* and *CAC* is no longer significant ($\beta = -0.004, p = 0.957$). The three-year lead results show a similar pattern, with the main effect of *Product Development* remaining positive and significant ($\beta = 0.188, p = 0.075$), and the interaction term remaining non-significant ($\beta = -0.045, p = 0.480$).

These findings suggest that while the direct effect of product development on future hiring persists, the moderating effect of CAC on this relationship diminishes over time. This could indicate that the mediocrity trap is primarily a short-term phenomenon, with its effects potentially dissipating over multiple years. The diminishing significance of the key interaction over time aligns with the dynamic nature of the B2B SaaS market. Firms that persist despite high CAC may eventually improve their product-market fit or exit, leading to a natural attenuation of the mediocrity trap effect in the long run. These results underscore the importance of timely product development and market feedback interpretation, as the consequences of falling into the mediocrity trap appear most pronounced in the near term. Future research could explore the specific mechanisms through which firms escape or adapt to this trap over time.

Table C.10: Sell Again: Two and Three Year Lead. Ordinary least squares (OLS) estimation. Standard errors (in parentheses) are clustered at the year level. *p*-values are shown in brackets.

	Sell Again	
	Two Year Lead	Three Year Lead
Ln Product Development	0.231 (0.057) [0.005]	0.188 (0.090) [0.075]
Ln CAC	0.024 (0.090) [0.799]	0.157 (0.112) [0.204]
Ln Product Development × Ln CAC	−0.004 (0.068) [0.957]	−0.045 (0.061) [0.480]
Ln Funding	1.751 (0.326) [0.001]	1.695 (0.380) [0.003]
Ln Total Employees	0.467 (0.249) [0.103]	−0.206 (0.225) [0.391]
Observations	8795	8098
R^2	0.513	0.565
FE: Year	X	X
FE: Firm	X	X

C.1.15 POST-HOC ANALYSIS

This section provides additional detail on the post hoc analysis briefly summarized in Section 3.6.1, which investigates the long-term implications of a firm's early resource allocation decisions. Specifically, I examine how early investments in sales versus product development—measured by the number of employees in each function during a firm's first two years—relate to downstream outcomes such as total customers, unicorn status, and shutdown risk. This section outlines the construction of the relevant variables and presents regression results to further explore the relationships discussed in the main text.

IMPACT ON CUSTOMER BASE I first examine the relationship between early sales and product development efforts and the long-term customer base. Table 3.5 shows a negative relationship between salespeople in the first two years and total customers in the long run ($\beta = -0.043, p = 0.076$). Product developers in the first two years also negatively correlate with total customers ($\beta = -0.161, p \approx 0.000$). However, the interaction term is positive and significant ($\beta = 0.022, p \approx 0.000$), suggesting that a balanced approach leads to better customer acquisition. A 100% increase in salespeople in the first two years associates with a 3.0% decrease in total customers. The positive interaction indicates that a 100% increase in both leads to a 1.5% increase in total customers.

IMPACT ON FIRM VALUATION I next examine the relationship between early resource allocation decisions and the likelihood of achieving “unicorn” status (a valuation of over \$1 billion). The results demonstrate a negative relationship between salespeople in the first two years and unicorn probability ($\beta = -0.014, p = 0.002$), as well as between product developers and unicorn status ($\beta = -0.011, p = 0.004$). The interaction term is positive and significant ($\beta = 0.004, p = 0.003$). A 100% increase in salespeople in the first two years associates with a 1.0 percentage point decrease in unicorn probability. Firms balancing both see a 0.3 percentage point increase in unicorn probability for a 100% increase in

both salespeople and product developers.

IMPACT ON FIRM SURVIVAL I then analyze the relationship between early resource allocation and firm survival. Table 3.5 indicates a significant negative interaction term, suggesting that balancing efforts mitigates failure risk, with a 100% increase in both salespeople and product developers associating with a 0.1 percentage point decrease in shutdown probability.

TIME TO SHUTDOWN Finally, I examine whether early resource allocation decisions affected the time to shutdown for firms that eventually failed. More failure is not, in and of itself, a bad thing; failure can provide meaningful learning. In fact, if a market does not exist, an entrepreneur's ideal outcome is to fail as quickly as possible so she can avoid wasting resources on perceived opportunities that do not exist in reality. However, the mediocrity trap may lead firms with lower-quality products to persist in the market, taking more time to fail compared to entering with a higher-quality product. Table 3.5 shows a negative relationship between product developers and time to shutdown ($\beta = -0.058, p = 0.159$). The interaction term is negative ($\beta = -0.024, p = 0.063$). A 100% increase in both salespeople and product developers in the first two years associates with a 1.7% decrease in time to shutdown.

In summary, this post hoc analysis reveals complex relationships between early resource allocation decisions and long-term outcomes. Over-emphasis on either sales or product development in the first two years appears to have negative implications, while a balanced approach leads to better outcomes in customer acquisition and high valuations. However, this balanced approach may also lead to quicker failure for unsuccessful firms, potentially indicating more efficient market exit.

Bibliography

- Abernathy, W. J. & Clark, K. B. (1985). Innovation: Mapping the winds of creative destruction. *Research Policy*, 14(1), 3–22.
- Adner, R. & Kapoor, R. (2016). Innovation ecosystems and the pace of substitution: Re-examining technology S-curves. *Strategic Management Journal*, 37(4), 625–648.
- Adner, R. & Levinthal, D. (2008). Doing versus seeing: acts of exploitation and perceptions of exploration. *Strategic Entrepreneurship Journal*, 2(1), 43–52.
- Adner, R. & Zemsky, P. (2006). A demand-based perspective on sustainable competitive advantage. *Strategic Management Journal*, 27(3), 215–239.
- Afuah, A. & Tucci, C. L. (2012). Crowdsourcing as a solution to distant search. *Academy of Management Review*, 37(3), 355–375.
- Agarwal, R., Bacco, F., Camuffo, A., Coali, A., Gambardella, A., Msangi, H., Sonka, S., Temu, A., Waized, B., & Wormald, A. (2025). Does a theory-of-value add value? Evidence from a randomized control trial with Tanzanian entrepreneurs. *Organization Science*.
- Agarwal, R. & Bayus, B. L. (2002). The market evolution and sales takeoff of product innovations. *Management Science*, 48(8), 1024–1041.
- Agarwal, R., Echambadi, R., Franco, A. M., & Sarkar, M. B. (2004). Knowledge transfer through inheritance: Spin-out generation, development, and survival. *Academy of Management Journal*, 47(4), 501–522.
- Aggarwal, V. A., Hsu, D. H., & Wu, A. (2020). Organizing knowledge production teams within firms for innovation. *Strategy Science*, 5(1), 1–16.
- Aghion, P., Bloom, N., Blundell, R., Griffith, R., & Howitt, P. (2005). Competition and innovation: An inverted-U relationship. *Quarterly Journal of Economics*, 120(2), 701–728.
- Agrawal, A., Gans, J. S., & Stern, S. (2021). Enabling entrepreneurial choice. *Management Science*, 67(9), 5510–5524.

- Ahuja, G., Coff, R. W., & Lee, P. M. (2005). Managerial foresight and attempted rent appropriation: Insider trading on knowledge of imminent breakthroughs. *Strategic Management Journal*, 26(9), 791–808.
- Ahuja, G. & Katila, R. (2004). Where do resources come from? The role of idiosyncratic situations. *Strategic management journal*, 25(8-9), 887–907.
- Alvarez, S. A. & Barney, J. B. (2007). Discovery and creation: Alternative theories of entrepreneurial action. *Strategic Entrepreneurship Journal*, 1(1), 11–26.
- Amit, R. & Zott, C. (2001). Value creation in E-business. *Strategic Management Journal*, 22(6–7), 493–520.
- Anderson, P. (1999). Perspective: Complexity theory and organization science. *Organization Science*, 10(3), 216–232.
- Artinger, S. & Powell, T. C. (2016). Entrepreneurial failure: Statistical and psychological explanations. *Strategic Management Journal*, 37(6), 1047–1064.
- Baker, T. & Nelson, R. E. (2005). Creating something from nothing: Resource construction through entrepreneurial bricolage. *Administrative Science Quarterly*, 50(3), 329–366.
- Baldwin, C. & Clark, K. (2000). *Design rules: The power of modularity*. Cambridge, MA: MIT Press.
- Baldwin, C. Y. & Clark, K. B. (2006). Between knowledge and the economy: Notes on the scientific study of designs. In B. Kahin & D. Foray (Eds.), *Advancing Knowledge and The Knowledge Economy*. MIT Press.
- Barnett, M. L. & Pollock, T. G. (2012). *The Oxford handbook of corporate reputation*. Oxford University Press.
- Barney, J. B. (1986). Strategic factor markets: Expectations, luck, and business strategy. *Management Science*, 32(10), 1231–1241.
- Barney, J. B. (1991). Firm resources and sustained competitive advantage. *Journal of Management*, 17(1), 99–120.
- Benzell, S. G., Hersh, J., & Van Alstyne, M. (2023). How APIs create growth by inverting the firm. *Management Science*.
- Bhat, N., Farias, V. F., Moallemi, C. C., & Sinha, D. (2020). Near-optimal A-B testing. *Management Science*, 66(10), 4477–4495.
- Brandenburger, A. M. & Stuart Jr., H. W. (1996). Value-based business strategy. *Journal of Economics & Management Strategy*, 5(1), 5–24.

- Brown, S. L. & Eisenhardt, K. M. (1997). The art of continuous change: Linking complexity theory and time-paced evolution in relentlessly shifting organizations. *Administrative science quarterly*, (pp. 1–34).
- Camerer, C. & Lovo, D. (1999). Overconfidence and excess entry: An experimental approach. *American Economic Review*, 89(1), 306–318.
- Camuffo, A., Cordova, A., Gambardella, A., & Spina, C. (2020). A scientific approach to entrepreneurial decision making: Evidence from a randomized control trial. *Management Science*, 66(2), 564–586.
- Carnabuci, G. & Operti, E. (2013). Where do firms' recombinant capabilities come from? Intra-organizational networks, knowledge, and firms' ability to innovate through technological recombination. *Strategic Management Journal*, 34(13), 1591–1613.
- Carreyrou, J. (2018). *Bad blood: Secrets and lies in a Silicon Valley startup*. New York: Alfred A. Knopf.
- Cassar, G. (2010). Are individuals entering self-employment overly optimistic? An empirical test of plans and projections on nascent entrepreneur expectations. *Strategic Management Journal*, 31(8), 822–840.
- Cassar, G. (2014). Industry and startup experience on entrepreneur forecast performance in new firms. *Journal of Business Venturing*, 29(1), 137–151.
- Chandler, A. D. (1966). *Strategy and structure: Chapters in the history of the American enterprise*. Cambridge, MA: Massachusetts Institute of Technology.
- Chen, J. S., Croson, D. C., Elfenbein, D. W., & Posen, H. E. (2019). The impact of learning and overconfidence on entrepreneurial entry and exit. *Organization Science*, 29(6), 989–1009.
- Chen, J. S., Elfenbein, D. W., Posen, H. E., & Wang, M. z. (2024). Programs of experimentation and pivoting for (overconfident) entrepreneurs. *Academy of Management Review*, 49(1), 80–106.
- Cho, H. & Clough, D. (2015). 'If at first you don't succeed': Learning and aspiration level adjustment in venture fundraising on Kickstarter. In *Academy of Management Annual Meeting*.
- Christensen, C. (1997). *The innovator's dilemma*. Boston: Harvard Business Review Press.
- Clough, D. R., Fang, T. P., Vissa, B., & Wu, A. (2018). Turning lead into gold: How do entrepreneurs mobilize resources to exploit opportunities? *Academy of Management Annals*, 13(1), 240–271.
- Clough, D. R., Fang, T. P., Vissa, B., & Wu, A. (2019). Turning lead into gold: How do entrepreneurs mobilize resources to exploit opportunities? *Academy of Management Annals*, 13(1), 240–271.

- Coase, R. H. (1937). The nature of the firm. *Economica*, 4(16), 386–405.
- Cohen, S. L., Bingham, C. B., & Hallen, B. L. (2019). The role of accelerator designs in mitigating bounded rationality in new ventures. *Administrative Science Quarterly*, 64(4), 810–854.
- Cohen, W. M. & Levinthal, D. A. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly*, 35(1), 128–152.
- Cozzolino, A. & Rothaermel, F. T. (2018). Discontinuities, competition, and cooperation: Coopetitive dynamics between incumbents and entrants. *Strategic Management Journal*, 39(12), 3053–3085.
- Csaszar, F. A. (2012). Organizational structure as a determinant of performance: Evidence from mutual funds. *Strategic Management Journal*, 33(6), 611–632.
- Csaszar, F. A. (2018). What makes a decision strategic? Strategic representations. *Strategy Science*, 3(4), 606–619.
- Csaszar, F. A. & Laureiro-Martínez, D. (2018). Individual and organizational antecedents of strategic foresight: A representational approach. *Strategy Science*, 3(3), 513–532.
- Cyert, R. M. & March, J. G. (1963). *A behavioral theory of the firm*. Englewood Cliffs, NJ: Prentice-Hall.
- Dahl, M. S. & Sorenson, O. (2012). Home sweet home: Entrepreneurs' location choices and the performance of their ventures. *Management Science*, 58(6), 1059–1071.
- Darr, E. D., Argote, L., & Epple, D. (1995). The acquisition, transfer, and depreciation of knowledge in service organizations: Productivity in franchises. *Management Science*, 41(11), 1750–1762.
- Denrell, J., Fang, C., & Levinthal, D. A. (2004). From T-mazes to labyrinths: Learning from model-based feedback. *Management Science*, 50(10), 1366–1378.
- DeSantola, A. & Gulati, R. (2017). Scaling: Organizing and growth in entrepreneurial ventures. *Academy of Management Annals*, 11(2), 640–668.
- Dierickx, I. & Cool, K. (1989). Asset stock accumulation and sustainability of competitive advantage. *Management Science*, 35(12), 1504–1511.
- Edmondson, A. C., Dillon, J. R., & Roloff, K. S. (2007). Three perspectives on team learning. *Academy of Management Annals*, 1(1), 269–314.
- Eggers, J. P. & Kaul, A. (2018). Motivation and ability? A behavioral perspective on the pursuit of radical invention in multi-technology incumbents. *Academy of Management Journal*, 61(1), 67–93.
- Eggers, J. P. & Song, L. (2015). Dealing with failure: Serial entrepreneurs and the costs of changing industries between ventures. *Academy of Management Journal*, 58(6), 1785–1803.

- Eisenhardt, K. M. & Bingham, C. B. (2017). Superior strategy in entrepreneurial settings: Thinking, doing, and the logic of opportunity. *Strategy Science*, 2(4), 246–257.
- Elfenbein, D. W. & Knott, A. M. (2015). Time to exit: Rational, behavioral, and organizational delays. *Strategic Management Journal*, 36(7), 957–975.
- Elfenbein, D. W., Knott, A. M., & Croson, R. (2017). Equity stakes and exit: An experimental approach to decomposing exit delay. *Strategic Management Journal*, 38(2), 278–299.
- Ethiraj, S. K. (2007). Allocation of inventive effort in complex product systems. *Strategic Management Journal*, 28(6), 563–584.
- Ethiraj, S. K., Kale, P., Krishnan, M. S., & Singh, J. V. (2005). Where do capabilities come from and how do they matter? A study in the software services industry. *Strategic Management Journal*, 26(1), 25–45.
- Ethiraj, S. K. & Levinthal, D. (2009). Hoping for A to Z while rewarding only A: Complex organizations and multiple goals. *Organization Science*, 20(1), 4–21.
- Ethiraj, S. K., Ramasubbu, N., & Krishnan, M. S. (2012). Does complexity deter customer-focus? *Strategic Management Journal*, 33(2), 137–161.
- Ewens, M., Rhodes-Kropf, M., & Strebulaev, I. A. (2016). Insider financing and venture capital returns.
- Felin, T., Gambardella, A., Stern, S., & Zenger, T. (2020). Lean startup and the business model: Experimentation revisited. *Long Range Planning*, 53(4), 101889.
- Felin, T., Kauffman, S., & Zenger, T. (2023). Resource origins and search. *Strategic Management Journal*, 44(6), 1514–1533.
- Felin, T. & Zenger, T. R. (2017). The theory-based view: Economic actors as theorists. *Strategy Science*, 2(4), 258–271.
- Fleming, L. (2001). Recombinant uncertainty in technological search. *Management Science*, 47(1), 117–132.
- Fleming, L. & Sorenson, O. (2001). Technology as a complex adaptive system: Evidence from patent data. *Research Policy*, 30(7), 1019–1039.
- Forbes, D. P. (2005). Are some entrepreneurs more overconfident than others? *Journal of Business Venturing*, 20(5), 623–640.
- Fortunato, S. & Barthelemy, M. (2007). Resolution limit in community detection. *Proceedings of the National Academy of Sciences*, 104(1), 36–41.

- Furr, N. & Ahlstrom, P. (2011). *Nail it then scale it: the entrepreneur's guide to creating and managing breakthrough innovation*, volume 658. Nisi Institute Lexington, KY.
- Gans, J. S., Stern, S., & Wu, J. (2019). Foundations of entrepreneurial strategy. *Strategic Management Journal*, 40(5), 736–756.
- Gary, M. S. & Wood, R. E. (2011). Mental models, decision rules, and performance heterogeneity. *Strategic Management Journal*, 32(6), 569–594.
- Gavetti, G. (2012). Toward a behavioral theory of strategy. *Organization Science*, 23(1), 267–285.
- Gavetti, G. & Menon, A. (2016). Evolution cum agency: Toward a model of strategic foresight. *Strategy Science*, 1(3), 207–233.
- Gompers, P., Kovner, A., Lerner, J., & Scharfstein, D. (2010). Performance persistence in entrepreneurship. *Journal of Financial Economics*, 96(1), 18–32.
- Grégoire, D. A. & Shepherd, D. A. (2012). Technology-market combinations and the identification of entrepreneurial opportunities: An investigation of the opportunity-individual nexus. *Academy of Management Journal*, 55(4), 753–785.
- Greve, H. R. (1998). Performance, aspirations, and risky organizational change. *Administrative Science Quarterly*, 43(1), 58–86.
- Greve, H. R. (2003). A behavioral theory of R&D expenditures and innovations: Evidence from shipbuilding. *Academy of Management Journal*, 46(6), 685–702.
- Gruber, M., Harhoff, D., & Hoisl, K. (2012). Knowledge recombination across technological boundaries: Scientists vs. engineers. *Management Science*, 59(4), 837–851.
- Gruber, M., MacMillan, I. C., & Thompson, J. D. (2008). Look before you leap: Market opportunity identification in emerging technology firms. *Management Science*, 54(9), 1652–1665.
- Hallen, B. L. & Pahnke, E. C. (2016). When do entrepreneurs accurately evaluate venture capital firms' track records? A bounded rationality perspective. *Academy of Management Journal*, 59(5), 1535–1560.
- Hambrick, D. C. & Crozier, L. M. (1985). Stumblers and stars in the management of rapid growth. *Journal of Business Venturing*, 1(1), 31–45.
- Hayward, M. L., Shepherd, D. A., & Griffin, D. (2006). A hubris theory of entrepreneurship. *Management Science*, 52(2), 160–172.
- Helfat, C. E. & Lieberman, M. B. (2002). The birth of capabilities: market entry and the importance of pre-history. *Industrial and Corporate Change*, 11(4), 725–760.

- Hitt, M. A., Bierman, L., Shimizu, K., & Kochhar, R. (2001). Direct and moderating effects of human capital on strategy and performance in professional service firms: A resource-based perspective. *Academy of Management Journal*, 44(1), 13–28.
- Hogarth, R. M. & Karelaia, N. (2012). Entrepreneurial success and failure: Confidence and fallible judgment. *Organization Science*, 23(6), 1733–1747.
- Hsu, D. H. (2007). Experienced entrepreneurial founders, organizational capital, and venture capital funding. *Research Policy*, 36(5), 722–741.
- Huang, L. (2017). The role of investor gut feel in managing complexity and extreme risk. *Academy of Management Journal*, 61(5), 1821–1847.
- Huber, G. P. (1991). Organizational learning: The contributing processes and the literatures. *Organization Science*, 2(1), 88–115.
- Hurst, E. & Pugsley, B. W. (2011). What do small businesses do? *National Bureau of Economic Research Working Paper Series*, No. 17041.
- Joseph, J. & Gaba, V. (2015). The fog of feedback: Ambiguity and firm responses to multiple aspiration levels. *Strategic Management Journal*, 36(13), 1960–1978.
- Kaplan, S. N., Sensoy, B. A., & Strömberg, P. (2009). Should investors bet on the jockey or the horse? Evidence from the evolution of firms from early business plans to public companies. *Journal of Finance*, 64(1), 75–115.
- Kapoor, R. & Wilde, D. (2020). Peering into a crystal ball: Foresight during periods of industry change. *Wharton School Working Paper*.
- Kapoor, R. & Wilde, D. (2023). Peering into a crystal ball: Forecasting behavior and industry foresight. *Strategic Management Journal*, 44(3), 704–736.
- Katila, R. & Shane, S. (2005). When does lack of resources make new firms innovative? *Academy of Management Journal*, 48(5), 814–829.
- Kauffman, S. A. (1995). *At home in the universe*. New York: Oxford University Press.
- Keijl, S., Gilsing, V., & Duysters, G. (2012). The trap between local and distant new combinations: Recombination and impact of inventions. *Academy of Management Proceedings*, 2012(10613).
- Keller, R. T. (2001). Cross-functional project groups in research and new product development: Diversity, communications, job stress, and outcomes. *Academy of Management Journal*, 44(3), 547–555.
- Kemper, S. (2003). *Code name Ginger: the story behind Segway and Dean Kamen's quest to invent a new world*. Harvard Business Press.

- Kerr, W. R., Nanda, R., & Rhodes-Kropf, M. (2014). Entrepreneurship as experimentation. *Journal of Economic Perspectives*, 28(3), 25–48.
- Keum, D. D. & Eggers, J. P. (2018). Setting the bar: The evaluative and allocative roles of organizational aspirations. *Organization Science*, 29(6), 1170–1186.
- Keynes, J. M. (1937). The general theory of employment. *Quarterly Journal of Economics*, 51(2), 209–223.
- Kickstarter PBC (2020). *Our rules*. <https://www.kickstarter.com/rules>.
- Kim, W. C. & Mauborgne, R. (2014). *Blue ocean strategy, expanded edition: How to create uncontested market space and make the competition irrelevant*. Harvard business review Press.
- Kneeland, M. K., Schilling, M. A., & Aharonson, B. S. (2020). Exploring uncharted territory: Knowledge search processes in the origination of outlier innovation. *Organization Science*, 31(3), 535–557.
- Knight, F. H. (1921). *Risk, uncertainty and profit*. Boston: Houghton Mifflin.
- Knudsen, T., Levinthal, D. A., & Winter, S. G. (2014). Hidden but in plain sight: The role of scale adjustment in industry dynamics. *Strategic Management Journal*, 35(11), 1569–1584.
- Kohli, A. K. & Jaworski, B. J. (1990). Market orientation: The construct, research propositions, and managerial implications. *Journal of Marketing*, 54(2), 1–18.
- Koning, R., Hasan, S., & Chatterji, A. (2022). Experimentation and Start-up Performance: Evidence from A/B Testing. *Management Science*, 68(9), 6434–6453.
- Koudstaal, M., Sloof, R., & Van Praag, M. (2016). Risk, uncertainty, and entrepreneurship: Evidence from a lab-in-the-field experiment. *Management Science*, 62(10), 2897–2915.
- Kroll, M., Walters, B. A., & Le, S. A. (2007). The impact of board composition and top management team ownership structure on post-IPO performance in young entrepreneurial firms. *Academy of Management Journal*, 50(5), 1198–1216.
- Leatherbee, M. & Katila, R. (2020). The lean startup method: Early-stage teams and hypothesis-based probing of business ideas. *Strategic Entrepreneurship Journal*, 14(4), 570–593.
- Lee, M. & Huang, L. (2018). Gender bias, social impact framing, and evaluation of entrepreneurial ventures. *Organization Science*, 29(1), 1–16.
- Lee, S. R. & Kim, J. D. (2024). When do startups scale? Large-scale evidence from job postings. *Strategic Management Journal*, 9(45), 1633–1669.
- Leiblein, M. J., Reuer, J. J., & Zenger, T. (2018). What makes a decision strategic? *Strategy Science*, 3(4), 558–573.

- Levine, R. & Rubinstein, Y. (2017). Smart and illicit: Who becomes an entrepreneur and do they earn more? *Quarterly Journal of Economics*, 132(2), 963–1018.
- Levinthal, D. A. (1997). Adaptation on rugged landscapes. *Management Science*, 43(7), 934–950.
- Levinthal, D. A. & March, J. G. (1993). The myopia of learning. *Strategic Management Journal*, 14(S2), 95–112.
- Levitt, S. D., List, J. A., & Syverson, C. (2013). Toward an understanding of learning by doing: Evidence from an automobile assembly plant. *Journal of Political Economy*, 121(4), 643–681.
- Li, E. & Martin, J. S. (2019). Capital formation and financial intermediation: The role of entrepreneur reputation formation. *Journal of Corporate Finance*, 59, 185–201.
- Lieberman, M. B. (1989). The learning curve, technology barriers to entry, and competitive survival in the chemical processing industries. *Strategic Management Journal*, 10(5), 431–447.
- March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science*, 2(1), 71–87.
- McDonald, R. M. & Eisenhardt, K. M. (2019). Parallel Play: Startups, Nascent Markets, and Effective Business-model Design. *Administrative Science Quarterly*, 65(2), 483–523.
- Minniti, M. & Bygrave, W. (2001). A dynamic model of entrepreneurial learning. *Entrepreneurship Theory and Practice*, 25(3), 5–16.
- Moeen, M. & Agarwal, R. (2017). Incubation of an industry: Heterogeneous knowledge bases and modes of value capture. *Strategic Management Journal*, 38(3), 566–587.
- Moeen, M., Agarwal, R., & Shah, S. K. (2020). Building Industries by Building Knowledge: Uncertainty Reduction over Industry Milestones. *Strategy Science*, 5(3), 218–244.
- Mollick, E. (2014). The dynamics of crowdfunding: An exploratory study. *Journal of Business Venturing*, 29(1), 1–16.
- Mollick, E. (2015). Delivery rates on Kickstarter. *SSRN Working Paper*.
- Mukherjee, S., Uzzi, B., Jones, B., & Stringer, M. (2016). A new method for identifying recombinations of existing knowledge associated with high-impact innovation. *Journal of Product Innovation Management*, 33(2), 224–236.
- Nanda, R. & Rhodes-Kropf, M. (2016). Financing risk and innovation. *Management Science*, 63(4), 901–918.
- Navis, C. & Glynn, M. A. (2010). How new market categories emerge: Temporal dynamics of legitimacy, identity, and entrepreneurship in satellite radio, 1990–2005. *Administrative Science Quarterly*, 55(3), 439–471.

- Nelson, R. & Winters, S. (1982). *An evolutionary theory of economic change*. Cambridge, MA: Harvard University Press.
- Nelson, R. R. & Winter, S. G. (2002). Evolutionary theorizing in economics. *Journal of Economic Perspectives*, 16(2), 23–46.
- Nicholls-Nixon, C. L. (2005). Rapid growth and high performance: The entrepreneur’s “impossible dream?”. *Academy of Management Perspectives*, 19(1), 77–89.
- Nickerson, J. & Argyres, N. (2018). Strategizing before strategic decision making. *Strategy Science*, 3(4), 592–605.
- Olson, E. M., Walker, O. C., & Ruekert, R. W. (1995). Organizing for effective new product development: The moderating role of product innovativeness. *Journal of Marketing*, 59(1), 48–62.
- Ott, T. E. & Eisenhardt, K. M. (2020). Decision weaving: Forming novel, complex strategy in entrepreneurial settings. *Strategic Management Journal*, 41(12), 2275–2314.
- Ott, T. E., Eisenhardt, K. M., & Bingham, C. B. (2017). Strategy formation in entrepreneurial settings: Past insights and future directions. *Strategic Entrepreneurship Journal*, 11(3), 306–325.
- Paik, Y. (2014). Serial entrepreneurs and venture survival: Evidence from U.S. venture-capital-financed semiconductor firms. *Strategic Entrepreneurship Journal*, 8(3), 254–268.
- Penrose, E. (1959). *The theory of the growth of the firm*. New York: John Wiley & Sons, Ltd.
- Peter, L. J. & Hull, R. (1969). *The Peter principle*. London: Souvenir Press.
- Peterson, A. & Wu, A. (2021). Entrepreneurial learning and strategic foresight. *Strategic Management Journal*, 42(13), 2357–2388.
- Politis, D. (2005). The process of entrepreneurial learning: A conceptual framework. *Entrepreneurship: Theory and Practice*, 29(4), 399–424.
- Pons, P. & Latapy, M. (2006). Computing communities in large networks using random walks. *Journal of Graph Algorithms and Applications*, 10(2), 191–218.
- Pontikes, E. G. (2012). Two sides of the same coin: How ambiguous classification affects multiple audiences’ evaluations. *Administrative Science Quarterly*, 57(1), 81–118.
- Pope, D. G., Price, J., & Wolfers, J. (2018). Awareness reduces racial bias. *Management Science*, 64(11), 4988–4995.
- Popper, K. (1969). *Conjectures and refutations: The growth of scientific knowledge*. Routledge Classics. London: Routledge & Kegan Paul, 3rd ed. revised edition.
- Porter, M. (1985). *Competitive advantage*. New York: Free Press.

- Porter, M. & Siggelkow, N. (2008). Contextuality within activity systems and sustainability of competitive advantage. *Academy of Management Perspectives*, 22(2), 34–56.
- Porter, M. E. (1980). *Competitive strategy: Techniques for analyzing industries and competitors*. New York: Free Press.
- Posen, H. E., Leiblein, M. J., & Chen, J. S. (2018). Toward a behavioral theory of real options: Noisy signals, bias, and learning. *Strategic Management Journal*, 39(4), 1112–1138.
- Rapping, L. (1965). Learning and World War II production functions. *Review of Economics and Statistics*, 47(1), 81–86.
- Raveendran, M., Silvestri, L., & Gulati, R. (2020). The role of interdependence in the microfoundations of organization design: Task, goal, and knowledge interdependence. *Academy of Management Annals*, 14(2), 828–868.
- Reis, E. (2011). *The lean startup*. New York: Crown Business.
- Ries, E. (2011). *The lean startup*. New York: Crown Business.
- Rinker, T. W. (2019). {sentimentr}: Calculate text polarity sentiment. *GitHub Repository*.
- Rivkin, J. W. (2000). Imitation of complex strategies. *Management Science*, 46(6), 824–844.
- Rosvall, M. & Bergstrom, C. T. (2008). Maps of random walks on complex networks reveal community structure. *Proceedings of the National Academy of Sciences*, 105(4), 1118–1123.
- Rothaermel, F. T. & Hess, A. M. (2007). Building dynamic capabilities: Innovation driven by individual-, firm-, and network-level effects. *Organization Science*, 18(6), 898–921.
- Scherer, F. M. (1967). Research and development resource allocation under rivalry. *Quarterly Journal of Economics*, 81(3), 359–394.
- Schumpeter, J. (1939). *Business cycles*. New York: McGraw-Hill.
- Simon, H. A. (1962). The architecture of complexity. *Proceedings of the American Philosophical Society*, 106(6), 467–482.
- Simon, H. A. (1969). *The sciences of the artificial*. Cambridge, MA: MIT Press.
- Simon, H. A. (1990). Invariants of human behavior. *Annual Review of Psychology*, 41(1), 1–20.
- Singh, J. & Fleming, L. (2009). 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. In L. Cumming & B. Staw (Eds.), *Research in organizational behavior* (pp. 231–266). New York: JAI Press.

- Sommer, S. C. & Loch, C. H. (2004). Selectionism and learning in projects with complexity and unforeseeable uncertainty. *Management Science*, 50(10), 1334–1347.
- Srikanth, K., Anand, J., & Stan, M. (2021). The origins of time compression diseconomies. *Strategic Management Journal*, 42(9), 1573–1599.
- Stroube, B. & Dushnitsky, G. (2023). Built with BuiltWith: A technical note. *Available at SSRN 4475309*.
- Stuart, R. W. & Abetti, P. A. (1990). Impact of entrepreneurial and management experience on early performance. *Journal of Business Venturing*, 5(3), 151–162.
- Stuart, T. E., Hoang, H., & Hybels, R. C. (1999). Interorganizational endorsements and the performance of entrepreneurial ventures. *Administrative Science Quarterly*, 44(2), 315–349.
- Suarez, F. F. & Lanzolla, G. (2007). The role of environmental dynamics in building a first mover advantage theory. *Academy of Management Review*, 32(2), 377–392.
- Talbot, F. B. (1982). Resource-constrained project scheduling with time-resource tradeoffs: The nonpreemptive case. *Management Science*, 28(10), 1197–1210.
- Tidhar, R. & Eisenhardt, K. M. (2020). Get rich or die trying... finding revenue model fit using machine learning and multiple cases. *Strategic Management Journal*, 41(7), 1245–1273.
- Toffler, A. (1985). *The adaptive corporation*. New York: McGraw-Hill.
- Townsend, D. M., Hunt, R. A., McMullen, J. S., & Sarasvathy, S. D. (2018a). Uncertainty, knowledge problems, and entrepreneurial action. *Academy of Management Annals*, 12(2), 659–687.
- Townsend, D. M., Hunt, R. A., McMullen, J. S., & Sarasvathy, S. D. (2018b). Uncertainty, knowledge problems, and entrepreneurial action. *Academy of Management Annals*, 12(2), 659–687.
- Traag, V. A., Waltman, L., & Van Eck, N. J. (2019). From Louvain to Leiden: Guaranteeing well-connected communities. *Scientific Reports*, 9(1), 5233.
- Ulrich, K., Eppinger, S., & Yang, M. (2020). *Product design and development*. New York: McGraw-Hill Education, 7th edition.
- Utterback, J. M. (1974). Innovation in industry and the diffusion of technology. *Science*, 183(4125), 620–626.
- Utterback, J. M. & Suárez, F. F. (1993). Innovation, competition, and industry structure. *Research Policy*, 22(1), 1–21.
- Uzzi, B., Mukherjee, S., Stringer, M., & Jones, B. (2013). Atypical combinations and scientific impact. *Science*, 342(6157), 468–472.

- Valentine, J., Novelli, E., & Agarwal, R. (2024). The theory-based view and strategic pivots: The effects of theorization and experimentation on the type and nature of pivots. *Strategy Science*, 9(4), 433–460.
- Van den Steen, E. (2016). A formal theory of strategy. *Management Science*, 63(8), 2616–2636.
- von Hippel, E. (1986). Lead users: A source of novel product concepts. *Management Science*, 32(7), 791–805.
- Wu, B. & Knott, A. M. (2006). Entrepreneurial risk and market entry. *Management Science*, 52(9), 1315–1330.
- Wuebker, R., Zenger, T., & Felin, T. (2023). The theory-based view: Entrepreneurial microfoundations, resources, and choices. *Strategic Management Journal*, 44(12), 2922–2949.
- Yao, D. A. (1988). Beyond the reach of the invisible hand: Impediments to economic activity, market failures, and profitability. *Strategic Management Journal*, 9(S1), 59–70.
- Zachariadis, M. & Ozcan, P. (2017). The API economy and digital transformation in financial services: The case of open banking.
- Zellweger, T. & Zenger, T. (2023). Entrepreneurs as scientists: A pragmatist approach to producing value out of uncertainty. *Academy of Management Review*, 48(3), 379–408.
- Zott, C. & Amit, R. (2007). Business model design and the performance of entrepreneurial firms. *Organization Science*, 18(2), 181–199.
- Zott, C. & Amit, R. (2008). The fit between product market strategy and business model: implications for firm performance. *Strategic Management Journal*, 29(1), 1–26.