Dynamic Models of Innovation and Learning

by

Alan Jaske

Department of Economics
Duke University

Date: ______________________________

Approved:

____
Attila Ambrus, Dissertation Advisor

____
James Anton

____
Gary Biglaiser

____
Leslie Marx

____
Curtis Taylor

Dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy
in the Department of Economics
in the Graduate School of
Duke University

2023

Abstract

Dynamic Models of Innovation and Learning

by

Alan Jaske

Department of Economics
Duke University

Date: ____________________________________

Approved:

____
Attila Ambrus, Dissertation Advisor

____
James Anton

____
Gary Biglaiser

____
Leslie Marx

____
Curtis Taylor

An abstract of a dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy
in the Department of Economics
in the Graduate School of
Duke University

2023

Abstract

This dissertation analyzes three dynamic models of innovation. The primary focus of the analysis of each model is to highlight the role of endogenous mechanisms—such as observational learning and product market competition—in determining firms' R&D incentives, the speed of innovation, and the efficiency of equilibrium investment.

In Chapter 2 of this dissertation, entitled "Learning and the Timing of New Technology Adoption," I develop a continuous-time model in which two asymmetric firms (incumbent and challenger) decide whether to adopt an innovation of uncertain quality; and if so, when. Early adoption is valuable, as it comes with the possibility of generating a first-mover advantage. However, early adoption comes at the price of greater uncertainty because the innovation's true quality can be learned only by observing market performance. In particular, firms learn through a process of market experimentation by which the innovation's true quality is revealed at a rate proportional to first-mover sales. Through this endogenous learning channel, the flow of information becomes inextricably linked to market structure, which I show has significant implications for equilibrium pricing incentives and the dynamics of innovation diffusion. In particular, my results provides an information-theoretic explanation for why established incumbents, despite possessing substantial innovative ability, are disproportionately slow to respond to disruptive innovation—commonly referred to as the Innovator's Dilemma.

In Chapter 3 of this dissertation, entitled "The Role of Market Structure in Competitive Experimentation," analyzes the strategic incentives of firms to trade-off investment

into new versus existing technology. Specifically, I develop a continuous-time model in which two firms choose how to allocate a stock of resources toward either R&D (to develop an possible innovation) or current production (to maximize short-term competitiveness). From a technical standpoint, this is the first model of strategic experimentation in which direct payoff externalities arise through endogenous product market interaction. In terms of empirical contribution, my model provides an explanation for (1) the sensitivity of results in empirical work on R&D competition to industry specification and (2) the seemingly paradoxical decisions of historically innovative companies (e.g. Polaroid, RCA, Xerox) to abandon R&D highly-promising new technologies. Finally, the model be used for policy analysis—namely, to evaluate the effect of horizontal mergers on innovation and welfare. Specifically, my results provide a micro-foundation for conceptual arguments made in recent high-profile merger cases (e.g. Dow/DuPont) based on an "innovation theory of harm," by which a statically pro-competitive merger may be viewed as dynamically anti-competitive due to its negative effects on dynamic innovation incentives.

In Chapter 4 of this dissertation, entitled "Leadership and the Value of Persistence," which is co-authored with James Anton and Dennis Yao, Co-authored work with James Anton (Duke Fuqua) and Dennis Yao (Harvard Business School) expands the scope of my existing research to include organizational incentives for innovation and the economics of leadership. In a project called "Leadership and the Value of Persistence," we seek to answer the following question: "What is the value of having a persistent leader in an organization?" To make this question precise, we develop and analyze a dynamic model of

organizational innovation. In the model, a leader faces a sequence of projects and wants to motivate self-interested managers to work on them so that they can succeed and generate value for the organization. The key strategic tension is that individual managers only internalize the portion of project returns they receive, but the leader internalizes the project's value to the entire organization. In equilibrium, leaders use persistence (i.e. the refusal to give up on failed projects) for several reasons. First, persistence helps overcome the problem of moral hazard in teams by creating dynamic incentives for self-interested managers to work on projects of low individual value but high organizational value. Second, we show that persistence can be used to credibly communicate private information about high value projects. Intuitively, a leader's decision to persist after failure provides good news to managers about the project's value which, in turn, motivates effort. However, persistence can also exacerbate the problem of moral hazard by creating strategic incentives for managers to "experiment with failure" by shirking on current projects in order to learn about the project's value through the leader's persistence decision.

Dedication

To all my friends, family, and loved ones, both past and present, who have encouraged and supported me throughout my life.

Contents

List of Figures

Acknowledgements

I want to acknowledge and thank my dissertation advisor and committee members, Attila Ambrus, Gary Biglaiser, James Anton, Leslie Marx, and Curtis Taylor, for providing outstanding guidance and support while writing this dissertation. Your insightful feedback has shaped my interests as an economist and my abilities as a researcher. I will never forget the incredible opportunities I received under your supervision—to learn new things, to meet new people, and to see the world.

Chapter 1

Introduction

In many sectors of the economy, firms have the opportunity to develop innovations or adopt technologies to improve products or service offerings. Economists have long been interested in the determinants of these innovation decisions, as they appear closely linked to economic growth (Comin and Hobijn, 2010), industrial competitiveness (Fagerberg et al., 2009), and welfare (Trajtenberg, 1989). Despite this longstanding interest, the economics literature has struggled to obtain a general theory to explain the plethora of empirical patterns observed in practice (Baker, 1995). These limitations are most noticeable in studies of firm characteristics, such as size or dominance, and their relationship to innovative performance.

In this dissertation, I contribute to the understanding of firm characteristics and innovative performance by analyzing three models of endogenous innovation. These models consider different aspect of the innovation process but are unified by three key themes. The first theme is that innovations are typically uncertain in terms of their desirability or feasibility. Firms therefore posses strong incentives to learn about the quality of new innovations via a process of strategic experimentation. The second theme is that successful innovation typically requires sacrificing short-term competitiveness. Consequently, the incentive of a firm to innovate or adopt new technologies critically depends both on the exogenous features of the innovation and the endogenous opportunities for alternative investment.

Chapter 2 considers the adoption and diffusion of an uncertain product innovation in a duopoly market. Firms are asymmetric and learn the innovation's true quality by observing first-mover sales, which gives rise to a rich set of equilibrium dynamics. First, the information generated by first-mover sales raises consumers' expectations about innovation quality but also accelerates second-mover adoption. To balance these dual considerations, first-movers engage in *penetration pricing*: setting a low initial price that sharply increases over time so as to reduce information flows to the second-mover. Second, initial quality advantage determines the speed of innovation diffusion through both an *opportunity-cost effect* as well as an *information effect*. My framework thus provides a unified explanation for why high-quality incumbents can be highly innovative as first-movers but also disproportionately slow to adopt innovations as second-movers. Third, incumbents and challengers endogenously select into distinct innovation timing strategies based on expectations about innovation quality. Incumbents are most likely to lead in developing highly-promising innovations, while challengers are most likely to lead in developing less-promising innovations. Furthermore, the endogenous selection that occurs in equilibrium favors *increasing dominance*.

In Chapter 3 of this dissertation, I develop a framework to examine the role of market structure on the incentives for competitive experimentation. In particular, this is the first paper to analyze competitive experimentation within an endogenous market context.

The model considers a patent race between two firms to develop a new technology. This new technology is known to be superior to firms' existing technology, but its feasibility is uncertain. Experimentation is required to successfully develop the new technology, but this comes at the expense of investment into existing technology and, hence, current profitability. Importantly, firms are direct competitors in both R&D and the product market. Thus, firms must confront a strategic trade-off between exploration of the new technology (to ensure future competitiveness) and exploitation of the existing technology (to ensure current competitiveness). In contrast to the existing literature, which typically assumes that firms are either potential entrants or exogenous competitors, my model allows firms to endogenously compete within a pre-existing market structure. This market structure is fully general, and my results are not tied to any particular specification. Instead, my analysis utilizies general properties of flow profits which are satisfied in most widely-used models of product market competition.

Finally, Chapter 4 considers intra-firm incentives for innovation and the importance of endogenous leadership. Co-authored with James Anton and Dennis Yao, this chapter considers a leader's decision whether to persist with an unsuccessful R&D project or to terminate the project in favor of a new project with an uncertain value. How does that decision affect the effort exerted by the manager assigned to the project? To study this question, we build and analyze an equilibrium, infinite-horizon model which embeds a search problem with an agency problem. We assess the policy value of a leader's persistence instrument under conditions of complete and incomplete information. Among other things, we find that persistence takes advantage of a manager’s incentive to gain access to future, potentially higher-payoff projects to induce effort on the current project. Furthermore, when the leader has superior information about the value of the current project, the manager may choose to delay effort to better take advantage of the information signal provided by the leader’s persistence choice.

Chapter 2

Learning and the Timing of New Technology Adoption

2.1 Introduction

A number of studies have found that incumbents are no less likely to introduce major product innovations compared to challengers (Blundell et al., 1999; Chandy and Tellis, 2000). At the same time, there is strong evidence that incumbents respond disproportionately slow to so-called "disruptive" innovation (Christensen, 1993; Henderson and Clark, 1990). This first body of work characterizes incumbents as neither incompetent nor unwilling to innovate as first-movers, but the second indicates that they are much less likely to do so in a timely fashion as second-movers. A primary challenge in developing a compelling model of innovation is to identify a mechanism that can explain such inconsistent patterns of innovative performance.

In this dissertation chapter, I argue that endogenous learning provides a unified explanation for the above findings and other widely-observed empirical phenomena. Specifically, I develop a continuous time model in which two asymmetric firms (incumbent and challenger) contemplate whether to adopt an innovation of uncertain quality; and if so, when. Early adoption is valuable, as it comes with the possibility of generating a first-mover advantage. However, early adoption comes at the price of uncertainty because the innovation's true quality can be learned only by observing market performance. In particular, firms learn through a process of *market experimentation* by which the innovation's true quality is revealed at a rate proportional to first-mover sales. Through this endogenous learning channel, the flow of information becomes inextricably linked to market structure, which I show has significant implications for equilibrium pricing incentives and the dynamics of innovation diffusion.

While endogenous learning is a theoretically-intuitive mechanism, there are in principle many reasons why a firm may prefer a specific timing of investment; for example, anticipated reductions in cost (as in Fudenberg and Tirole, 1985; Reinganum, 1983) or improvements in quality (as in Dutta et al., 1995; Riordan and Salant, 1994). Therefore, to justify my focus on learning as a central determinant of innovation diffusion, we should understand the extent to which it actually motivates the timing of firm investment.

To answer exactly this question, Klingebiel and Joseph (2016) examine the German mobile phone industry between 2004 and 2008, when the introduction of new hardware features was at an all-time high. Industry representatives in their sample regularly cited the trade-off between learning and strategic preemption as a *primary motivation* for the early or late adoption of new hardware features. One firm known to be early adopter described their motivations as follows: *"[Our firm] now tries to be an innovator . . . we have emphasized speed [and risk taking]"* (p. 1009). Another firm described their motivations for late adoption as follows: *"The longer we wait, the easier it is for us to gauge success. . . . [W]e kept monitoring forecasts while Samsung, Sony Ericsson, and a few others tried their luck* *at development. And the picture became clearer"* (p. 1009). The findings of Klingebiel and Joseph (2016) clearly indicate that firms' investment timing decisions were not driven by the inability to access key inputs or by anticipated changes in future cost or quality; rather, it was a strategic choice to either (a) capitalize on speed or (b) wait to learn from the success (or failure) of others.

Proceeding via backward induction, I first analyze the impact of endogenous learning on second-mover incentives. In the model, second-movers update their beliefs about the profitability of investment by observing first-mover sales, which generate confirmatory signals of the innovation's true quality. A bad signal confirms the innovation to be low quality and benefits the second-mover because it can avoid making an unprofitable investment. A good signal confirms the innovation to be high quality and hurts the second-mover because it must compete as a late adopter of the innovation. In equilibrium, second-mover incentives are intuitively characterized by a unique threshold belief above which adoption is optimal. The comparative statics of this threshold are determined by market fundamentals and the returns to quality. Equilibrium flow profits are convex and thus exhibit increasing returns to quality, which causes a firm with high initial quality to experience a high opportunity cost of waiting compared to its rival.

The above comparative static – the *cost-of-waiting effect* – implies that an incumbent with high initial quality will choose a lower second-mover threshold than a challenger with low initial quality. When learning is exogenous (i.e. independent of first-mover sales), the cost-of-waiting effect implies that the incumbent should, all else equal, be a faster second-mover than the challenger. But when learning is endogenous, second-mover adoption also depends on the speed of learning which is endogenously determined by first-mover sales. All else equal, the incumbent attracts more consumers as a second-mover in equilibrium and thus learns at a slower rate than the challenger. My first main result, Theorem 1, provides conditions under which this countervailing force – the *information effect* – unambiguously dominates the cost-of-waiting effect. As a result, the incumbent will appear to adopt innovations at a disproportionately slow rate as a second-mover compared to the challenger.

When learning is endogenous, the possibility of second-mover adoption significantly alters strategic pricing incentives. As in Bergemann and Välimäki (1997), both firms receive a positive value of information from first-mover sales. Thus, joint-learning incentives relax the intensity of price competition and lead to higher absolute prices compared to the case of exogenous learning. As for relative prices, Proposition 3 shows that first-movers engage in form of *penetration pricing* to defend their monopoly over the innovation. Specifically, the first-mover sets a low initial price so as to generate large amounts of information and raise beliefs about the innovation's quality. But as beliefs increase, the first-mover sharply raises its price so as to generate less information and, thus, delay second-mover adoption.

Turning to the question of initial adoption, I show that first-mover incentives account for both market structure and learning considerations. First, a high-quality first-mover experiences greater returns from adopting a high-quality innovation but also greater losses from adopting a low-quality innovation; I call this the *own-profit effect*. Second, a high-quality first-mover expects that its rival (due to the opportunity-cost effect) will choose a higher second-mover threshold; I call this the *rival-waiting effect*. Due to the convexity of flow profit, the own-profit effect lowers first-mover incentives, but the rival-waiting effect increases first-mover incentives. Under intuitive assumptions, Theorem 2 establishes that the own-profit effect dominates. Thus, initial adoption is characterized by a simple dichotomy: incumbents lead in developing highly-promising (i.e. safe) innovations, while challengers lead in developing less-promising (i.e. riskier) innovations.

The speed at which innovations diffuse in equilibrium may be socially excessive or insufficient compared to the social optimum. Furthermore, welfare gains would be realized if firms were to sequentially adopt innovations over wider range of initial beliefs than occurs in equilibrium. The intuition behind this result is that first-movers do not fully internalize the value of the information that first-mover sales generate for consumers and second-movers. Thus, firms are too reluctant to lead as first-movers but also too reluctant to wait as second-movers. A merger to monopoly can produce welfare gains by increasing the efficiency of adoption timing. Specifically, I identify conditions under which (i) the range of beliefs over which a monopolist chooses sequential adoption is a strict improvement over duopoly, and (ii) the speed at which innovations diffuse under monopoly is a strict improvement over duopoly for a wide range of beliefs. At the heart of these welfare gains is that a monopolist's enhanced market power allows it to partially internalize business-stealing effects and information externalities.

The remainder of this dissertation chapter is organized as follows. Section 2.2 describes the related literature. Section 2.3 introduces my basic framework. Section 2.4 characterizes the implications of endogenous learning for price competition and equilibrium adoption. Section 3.5 analyzes the social efficiency of equilibrium incentives and considers the welfare effects of a merger to monopoly. Section 3.7 concludes with a discussion of possible extensions and ideas for future work. All proofs and related calculations are located in the Appendix.

2.2 Related Literature

This dissertation chapter contributes to the literature on new technology adoption and the timing of innovation by highlighting the central role of endogenous learning in the dynamics of innovation diffusion. In particular, my framework shows that endogenous learning can explain several widely-observed empirical phenomena, all while remaining sufficiently tractable to allow for normative analysis of market structure, competition policy, and innovation policy.

As in the existing literature, I view technology adoption as an inherently dynamic process. Seminal works include Reinganum (1981) and Fudenberg and Tirole (1985) who demonstrate how strategic timing considerations can explain the gradual diffusion of innovation across firms. Each model considers a rich continuous-time environment but focuses on (i) ex-ante symmetric firms and (ii) an innovation of known profitability. By comparison, I give primary focus to the case where firms are *asymmetric* and innovation profitability is *uncertain*.

Early studies of the role of uncertainty in new technology adoption include Mamer and McCardle (1987) and Jensen (1982), who focus on symmetric duopoly, and Reinganum (1983), who focuses on an asymmetric (but static) Cournot oligopoly. More recent studies have incorporated additional elements of uncertainty, such as information spillovers (e.g. Décamps and Mariotti, 2004; Hoppe, 2000), social learning (e.g. Frick and Ishii, 2015; Wagner, 2018), and reputations (e.g. Hendricks, 1992). Each of these studies assume that players are ex-ante identical. By contrast, I give primary focus to the role of firm asymmetry. In a closely related work, Schivardi and Schneider (2008) consider an environment in which innovation quality is uncertain and firms are asymmetric. Through a series of numerical analyses, they show that a high-quality incumbent is *more likely* to adopt a uncertain new technology rather than less likely. Given their focus on exogenous learning, this simulation result is consistent with the cost-of-waiting effect in my framework.

The central feature of my framework is that the speed of learning and, hence, innovation diffusion is endogenously determined by first-mover sales. This assumption captures the idea that uncertainty regarding an innovation's true quality takes time to be resolved, as is often the case when there is two-sided learning in markets (e.g. Aghion et al., 1991; Bergemann and Välimäki, 1997; Judd and Riordan, 1994). Furthermore, I assume that firms observe confirmatory signals of the innovation's true quality and use the absence of signals to update beliefs in real-time. My framework therefore relates to others in the literature on strategic experimentation (e.g. Bolton and Harris, 1999; Keller and Rady, 2015; Keller et al., 2005).

The technical features of my model build off Bergemann and Välimäki (1997). In their model, an incumbent competes with an entrant of unknown quality. Both the incumbent and entrant learn by observing the entrant's performance in the market, and prices continuously adjust as belief-updating occurs. Despite these similarities, the focuses of our analyses are fundamentally different. The focus of Bergemann and Välimäki (1997) is price experimentation and learning dynamics when a new product becomes available to consumers. My focus is the adoption and diffusion of an innovation when it becomes available to firms. By making these adoption decisions explicit, my framework extends the scope of Bergemann and Valimaki's analysis by allowing firms to compete along two dimensions: *price* and *product design*. As a result, my framework uncovers new dynamic considerations that can be used to investigate a broad set of questions related to market structure and welfare.

2.3 Model

Two firms, I (the incumbent) and C (the challenger), compete in continuous time over an infinite horizon. Each has zero marginal cost, and is summarized by two exogenous characteristics: (i) quality $q_i \geq 0$, and (ii) horizontal features $\theta_i \in [0,1]$. For the sake of interpretation, I let the incumbent have higher initial quality than the challenger; that is, $q_I > q_C$. As in Bergemann and Välimäki (1997), the incumbent and challenger are located at opposite ends of a Hotelling linear city; it is without loss to set $\theta_I = 0$ and $\theta_C = 1$.

Consumers are short-lived and have quasi-linear preferences over quality, horizontal features, and wealth. Consumers have uniform preferences over quality but have idiosyncratic preferences over horizontal features. In particular, each consumer is indexed by a taste parameter $\theta \in [0,1]$ that determines the utility from consumption of Firm i's product:

$$u_\theta(q_i;i) = \begin{cases} q_i + h(1-\theta) & \text{if } i = 1, \\ \\ q_i + h\theta & \text{if } i = 2, \end{cases} \tag{1}$$

where $h > 0$ parameterizes the extent to which firms are horizontally differentiated. Prices subtract linearly from consumers' utilities; therefore, the net utility from consuming Firm i's product at the price p_i is simply $u_\theta(q_i;i) - p_i$.

The game begins at $t = 0$ when an exogenously-developed product innovation is discovered made available for firms to adopt. The quality of the innovation, denoted $\tilde{q}$, measures how much it contributes a firm's quality once adopted. The true value of $\tilde{q}$ is *uncertain* at time 0, but firms know it is either high ($\tilde{q} = q_H$) or low ($\tilde{q} = q_L$), where $q_H > 0 > q_L$.[1] Thus, a high-quality innovation ($\tilde{q} = q_H$) increases own-demand, but a low-quality innovation ($\tilde{q} = q_L$) decreases own-demand. The cost of adopting the innovation is normalized to zero, but reversal costs are assumed to be sufficiently high to render adoption practically irreversible.

My analysis can be done while remaining completely agnostic about the relative magnitudes of q_1, q_2, q_H, q_L, and h. The price of such generality, however, is a multitude of special cases. For the sake of convenience, I maintain the following assumption:

---

[1] Here, I make the implicit assumption that $\tilde{q}$ is the same for each firm; that is, the innovation's true quality is independent of idiosyncratic differences between firms. This appropriateness of this assumption will depend on context. In the case of hardware innovations, such as those examined by Klingebiel and Joseph (2016), the notion of "quality" is often determined by functionality or cost, and thus is most likely highly-correlated across firms.

Assumption 1 (Contestability). $|q_i - q_j| < h + \max\{q_H, -q_L\}$.

Assumption 1 is necessary and sufficient to guarantee that (i) the market is always covered, and (ii) that each firm always receives a positive market share. This simplifies the presentation of my analysis by ensuring that best-responses are given by the usual first-order conditions.

Firms and consumers possess symmetric initial beliefs that $\tilde{q} = q_H$ with probability $\alpha_0 \in (0, 1)$. Following Jensen (1982), I assume that agents learn about the innovation's quality from the arrival of public signals over time. The information provided in these signals can be categorized either as "good news" or "bad news." Hence, the available information at time t can be summarized a counting process (N_t^G, N_t^B) that tracks the arrival of good news and bad news over time.

In Jensen's original model, and in the follow-up literature it inspired, the arrival rate of news and, therefore, the speed of learning is taken to be exogenous. This is a strong assumption, which I relax in a similar manner to Bergemann and Välimäki (1997), Bonatti (2011), and Peitz et al. (2017) by allowing the arrival of news to depend on two factors: (i) the innovation's true quality, $\tilde{q}$, and (ii) the volume of sales (of the innovation) at time t, denoted n_t. In particular, I assume that a high (low) quality innovation generates good (bad) news at Poisson rate $\lambda_G n_t$ ($\lambda_B n_t$), where $\lambda_G > 0$ ($\lambda_B > 0$). Note that only high-quality innovations can generate good news, and only low-quality innovations can generate bad news. Thus, the arrival of news (good or bad) conclusively reveals the innovation's true quality, as is the case in Keller et al. (2005) (for good news only) and Keller and Rady

(2015) (for bad news only).

Posterior beliefs about innovation quality at time t are denoted α_t. An application of Bayes' rule gives the following expression for α_t conditional on no news:

$$\alpha_t = \frac{\alpha_0 e^{-\lambda_G \int_0^t n_s ds}}{\alpha_0 e^{-\lambda_G \int_0^t n_s ds} + (1-\alpha_0) e^{-\lambda_B \int_0^t n_s ds}}. \tag{2}$$

To compute the law of motion for α_t, differentiate (33) with respect to t and then simplify:

$$d\alpha_t = (\lambda_B - \lambda_G) n_t \alpha_t (1-\alpha_t) dt. \tag{3}$$

Observe that α_t is increasing in t (i.e. "no news is good news") whenever $\lambda_B - \lambda_G$ is positive, and that α_t is decreasing in t (i.e. "no news is bad news") whenever $\lambda_B - \lambda_G$ is negative. I obtain similar qualitative results for each case; however, the former generates a richer set of equilibrium dynamics. For the sake of exposition, I maintain the following assumption:

Assumption 2 (No News is Good News)**.** $\lambda_B > \lambda_G$.

A large body of research across economics, marketing, and psychology indicates that strategic timing advantages are ubiquitous (Gilbert and Birnbaum-More, 1996; Lieberman and Montgomery, 1988). To highlight the strategic tension between learning and the timing of adoption, I assume throughout my analysis that "early" adoption confers a strict first-mover advantage.

To formalize the distinction between "early" and "late" adoption, it is helpful to introduce the following terminology. If Firm i adopts the innovation *before* the arrival of good news, then I say that Firm i is an *early adopter* of the innovation; and if Firm i adopts the innovation *after* the arrival of good news, then I say that Firm i is a *late adopter* of the innovation. Early adopters benefit from a timing advantage over late adopters because they are able to, for example, secure access to high-quality inputs, establish product standards, or cultivate brand loyalty. Formally, I suppose that a firm's post-adoption quality as an early adopter is $q_i + \tilde{q}$, but its post-adoption quality as a late adopter is $q_i + \tilde{q} - x$, where $x \in [0, q_H]$ is a reduced-form measure of the extent to which there is a first-mover advantage.[2]

My solution concept is Markov-perfect equilibrium (MPE). In any MPE, Firm i's strategy is a triple $(p_i(\alpha), \phi_i(\alpha), \tau_i(\alpha))$. The first component, $p_i(\alpha)$, is the price Firm i charges given beliefs α and, implicitly, the stage of the game. The second component, $\phi_i(\alpha) \in [0, \infty]$, is the rate at which Firm i adopts the innovation given that Firm j has not yet adopted. Here, $\phi_i(\alpha) = \infty$ indicates that Firm i immediately adopts the innovation at beliefs α. The third component, $\tau_i(\alpha)$, is a stopping time that defines when Firm i adopts as a second-mover. Regarding the timing of actions, I adopt the convention that prices at time t are decided *after* adoption decisions are made. Finally, the definition of MPE is standard: Firm i's choice of (p_i, ϕ_i, τ_i) should maximize its expected present-discounted profit given (p_j, ϕ_j, τ_j) in every possible state, and vice versa for Firm j.

---

[2] Observe that Firm i benefits from a first-mover advantage if and only if Firm j is both a second-mover *and* a late-adopter. This subtle but important distinction is motivated by empirical work on market leadership – specifically Golder and Tellis (1993) – who show that pioneers of new markets or technologies rarely outperform so-called "fast second-movers."

2.4 Equilibrium Analysis

2.4.1 Price Competition

Begin by taking the adoption strategies (ϕ_1, τ_1) and (ϕ_2, τ_2) as given, and let $T_i \in [0, \infty]$ denote the realized calendar date at which Firm i adopts the innovation. Here, $T_i < \infty$ indicates that Firm i adopts the innovation in finite time, while $T_i = \infty$ indicates that Firm i never adopts the innovation. Using this notation, we can divide the game into three stages:

- Stage 1 (Neither firm has adopted): $t < \min\{T_1, T_2\}$,
- Stage 2 (One firm has adopted): $\min\{T_1, T_2\} \leq t < \max\{T_1, T_2\}$,
- Stage 3 (Both firms have adopted): $t \geq \max\{T_1, T_2\}$.

In Stages 1 and 3, the quality differential between Firm i and Firm j is constant; therefore, prices, market shares, and flow profits in MPE are identical to those in a "static" equilibrium where firms ignore dynamic considerations.

Lemma 1 (Static Equilibrium)**.** *Let $\Delta_{i,t} = q_{i,t} - q_{j,t}$ denote Firm i's quality advantage over Firm j at time t. Then Firm i's price, market share, and profit in the static equilibrium are*

$$p(\Delta_{i,t}) = h + \frac{1}{3}\Delta_{i,t}, \quad n(\Delta_{i,t}) = \frac{h + \frac{1}{3}\Delta_{i,t}}{2h}, \quad \textit{and} \quad \Pi(\Delta_{i,t}) = \frac{(h + \frac{1}{3}\Delta_{i,t})^2}{2h}.$$

In Stage 2, the quality differential between Firm i and Firm j depends non-trivially on current beliefs. Therefore, MPE prices necessarily involve dynamic considerations. To characterize these dynamics, suppose that Firm j is the first-mover in Stage 2. Given prices (p_j, p_i) and current beliefs $\alpha \in (0,1)$, the respective market shares that Firm i and Firm j receive are

$$n_i^F(\alpha) = \frac{h + \Delta_i - q(\alpha) + p_j - p_i}{2h} \quad \text{and} \quad n_j^L(\alpha) = \frac{h + \Delta_j + q(\alpha) + p_i - p_j}{2h}, \tag{4}$$

where $q(\alpha) = \alpha q_H + (1-\alpha) q_L$ denotes the innovation's current expected quality.

Let $V_i^F(\alpha)$ and $V_j^L(\alpha)$ denote each firm's respective value in Stage 2, and let $\lambda(\alpha) = \alpha\lambda_G + (1-\alpha)\lambda_B$ denote the expected per-sale arrival rate of news. By the Principle of Optimality, $V_i^F(\alpha)$ solves the equation

$$\begin{aligned} \rho V_i^F(\alpha) = \max_{p_i} \Big\{ & p_i n_i^F(\alpha) + \lambda(\alpha) n_j^L(\alpha) \big[\Omega_i^F(\alpha) - V_i^F(\alpha) \big] \\ & + (\lambda_B - \lambda_G) n_j^L(\alpha) \alpha (1-\alpha) (V_i^F)'(\alpha) \Big\}, \end{aligned} \tag{5}$$

where

$$\Omega_i^F(\alpha) = \frac{\alpha\lambda_G}{\lambda(\alpha)} \left(\frac{\Pi(\Delta_i - x)}{\rho} \right) + \frac{(1-\alpha)\lambda_B}{\lambda(\alpha)} \left(\frac{\Pi(\Delta_i - q_L)}{\rho} \right)$$

is Firm i's expected value after the arrival of news. Likewise, $V_j^L(\alpha)$ solves the equation

$$\begin{aligned} \rho V_j^L(\alpha) = \max_{p_j} \Big\{ & p_j n_j^L(\alpha) + \lambda(\alpha) n_j^L(\alpha) \big[\Omega_i^F(\alpha) - V_i^L(\alpha) \big] \\ & + (\lambda_B - \lambda_G) n_j^L(\alpha) \alpha (1-\alpha) (V_j^L)'(\alpha) \Big\}, \end{aligned} \tag{6}$$

where

$$\Omega_j^L(\alpha) = \frac{\alpha\lambda_G}{\lambda(\alpha)}\left(\frac{\Pi(\Delta_j + x)}{\rho}\right) + \frac{(1-\alpha)\lambda_B}{\rho}\left(\frac{\Pi(\Delta_j + q_L)}{\rho}\right).$$

Equations (5) and (6) decompose each firm's discounted value into three terms; in order, they are (i) flow profit, (ii) the expected gains/losses from news arrival, and (iii) the effect of belief-updating over time. While these decompositions are intuitive, the values functions they describe are highly complex. To maintain analytic tractability, I follow Bergemann and Välimäki (1997), Bolton and Harris (2000), and Keller and Rady (2020) and henceforth assume that firms do not discount future payoffs; instead, they use the strong long-run average criterion.

Definition 1 (Long-Run and Strong Long-Run Average)**.** *Let* $\{\Pi_t\}_{t\geq 0}$ *denote a (possibly random) stream of flow profit. The long-run average of* $\{\Pi_t\}_{t\geq 0}$ *starting at time* t *is*

$$\nu_t = \mathbb{E}_t\left[\lim_{T\to\infty}\frac{1}{T}\int_t^T \Pi_s ds\right].$$

Provided that $\{\nu_t\}_{t\geq 0}$ *is well-defined, the strong long-run average of* $\{\Pi_t\}_{t\geq 0}$ *starting at time* t *is*

$$\mathcal{V}_t = \mathbb{E}_t\left[\int_t^\infty [\Pi_s - \nu_s]ds\right].$$

In words, the strong long-run average computes the lifetime sum of expected profits net the long-run average at each instant of time. Specializing Definition 1 to the present setting, the long-run average profits of Firm i and Firm j are

$$v_i^F(\alpha) = \Pi(\Delta_i) + \theta_G(\alpha;\tau_i)[\Pi(\Delta_i - x) - \Pi(\Delta_i)] + \theta_B(\alpha;\tau_i)[\Pi(\Delta_i - q_L) - \Pi(\Delta_i)] \quad (7)$$

and

$$v_j^L(\alpha) = \Pi(\Delta_j) + \theta_G(\alpha;\tau_i)\big[\Pi(\Delta_j + x) - \Pi(\Delta_j)\big] + \theta_B(\alpha;\tau_i)\big[\Pi(\Delta_j + q_L) - \Pi(\Delta_j)\big], \quad (8)$$

where $\theta_G(\alpha;\tau_i)$ and $\theta_B(\alpha;\tau_i)$ denote the probabilities that good news and bad news, respectively, arrive before Firm i adopts the innovation as a second-mover. (Note: Explicit solutions for $\theta_G(\alpha;\tau_i)$ and $\theta_B(\alpha;\tau_j)$ are derived in the next section. For now, I take them as given.)

Using (7) and (8), we can then define the strong long-run average profits of Firm i and Firm j within the context of a particular MPE:

$$\mathcal{V}_i^F(\alpha) \equiv \mathbb{E}_0\left\{\int_0^\infty \big[p_i^F(\alpha_t)n_i^F(\alpha_t) - v_i^F(\alpha_t)\big]dt \;\middle|\; \alpha_0 = \alpha\right\}, \quad (9)$$

$$\mathcal{V}_j^L(\alpha) \equiv \mathbb{E}_0\left\{\int_0^\infty \big[p_j^L(\alpha_t)n_j^L(\alpha_t) - v_j^L(\alpha_t)\big]dt \;\middle|\; \alpha_0 = \alpha\right\}. \quad (10)$$

As in Keller and Rady (2020), I say that a pricing strategy is *reasonable* if it almost-surely converges to the full-information (i.e. static) MPE in finite time. Provided that $n_j^L(\alpha) > 0$, it is (almost-surely) guaranteed that news will arrive in finite time. Therefore, it is without loss of generality to restrict attention to reasonable pricing strategies.

The benefit of observing that all MPE pricing strategies are reasonable is two-fold. First, it guarantees that $\mathcal{V}_i^F(\alpha)$ and $\mathcal{V}_j^L(\alpha)$ are finite and satisfy the Principle of Optimality. Second, it allows us to derive analytically tractable versions of (5) and (6). To do this, let

$$\mathcal{G}(\alpha, V) = \alpha\lambda_G[V(1) - V(\alpha)] + (1-\alpha)\lambda_B[V(0) - V(\alpha)] + (\lambda_B - \lambda_G)\alpha(1-\alpha)V'(\alpha)$$

denote the infinitesimal generator of our Poisson learning process. Then Firm i's best-response program may be written as follows:

$$\begin{aligned} 0 &= \max_{p_i}\left\{p_i n_i^F(\alpha) - \nu_i^F(\alpha) + n_j^L(\alpha)\mathcal{G}(\alpha, \mathcal{V}_i^F(\alpha))\right\} \\ \iff 0 &= \max_{p_i}\left\{\frac{p_i n_i^F(\alpha) - \nu_i^F(\alpha)}{n_j^L(\alpha)}\right\} + \mathcal{G}(\alpha, \mathcal{V}_i^F(\alpha)). \end{aligned} \tag{11}$$

Likewise, Firm j's best-response program may be written as

$$\begin{aligned} 0 &= \max_{p_j}\left\{p_j n_j^L(\alpha) - \nu_j^L(\alpha) + n_j^L(\alpha)\mathcal{G}(\alpha, \mathcal{V}_j^L(\alpha))\right\} \\ \iff 0 &= \max_{p_j}\left\{\frac{p_j n_j^L(\alpha) - \nu_j^L(\alpha)}{n_j^L(\alpha)}\right\} + \mathcal{G}(\alpha, \mathcal{V}_j^L(\alpha)). \end{aligned} \tag{12}$$

In words, the strong long-run average criterion preserves analytic tractability by eliminating the continuation values $\mathcal{G}(\alpha, V_j^L)$ and $\mathcal{G}(\alpha, V_j^L)$ from each firm's best-response. Instead, dynamic considerations are reflected in the difference between current and long-run average profit. Solving (11) and (12) for $p_j^L(\alpha)$ and $p_i^L(\alpha)$ yields the following solution.

Proposition 1 (Equilibrium Prices)**.** *In Stage 2, there is a unique MPE in pricing strategies.*

For all $t < \tau_i$, the first-mover (Firm j) charges price

$$p_j^L(\alpha_t) = h + \frac{1}{3}\big(\Delta_j + q(\alpha_t)\big) - \frac{2}{3}\left(\delta(\alpha_t;\tau_i) - q_0(\alpha_t)\right), \tag{13}$$

where $\delta(\alpha_t;\tau_i) \equiv \theta_G(\alpha_t;\tau_i)x + \theta_B(\alpha_t;\tau_i)q_L$; the second-mover (Firm i) charges price

$$p_i^F(\alpha_t) = \frac{2}{3}\left(\Delta_i - \delta(\alpha_t;\tau_i)\right) + \sqrt{2h\nu_j^L(\alpha_t)}; \tag{14}$$

and the first-mover (Firm j) receives market share

$$n_j^L(\alpha_t) = \sqrt{\frac{\nu_j^L(\alpha_t)}{2h}}. \tag{15}$$

To provide intuition for Proposition 1, it is helpful to compare $p_j^L(\alpha_t)$ and $p_i^F(\alpha_t)$ to the static equilibrium prices, namely

$$\bar{p}_j^L(\alpha_t) \equiv p(\Delta_j + q(\alpha_t)) = h + \frac{1}{3}\big(\Delta_j + q(\alpha_t)\big)$$

and

$$\bar{p}_i^F(\alpha_t) \equiv p(\Delta_i - q(\alpha_t)) = h + \frac{1}{3}\left(\Delta_i - q(\alpha_t)\right).$$

By inspecting the solution for $p_j^L(\alpha_t)$, observe that

$$p_j^L(\alpha_t) = \bar{p}_j^L(\alpha_t) + \frac{2}{3}\left[\delta(\alpha_t;\tau_i) - q(\alpha_t)\right]. \tag{16}$$

Thus, Firm j's MPE price may be above or below the static equilibrium price, depending on whether its current quality advantage over Firm i (i.e. $\Delta_j + q(\alpha_t)$) is higher or lower than its long-run quality advantage (i.e. $\Delta_j + \delta(\alpha_t;\tau_i)$). If $q(\alpha_t) > \delta(\alpha_t;\tau_i)$, then Firm j expects its quality advantage to decrease in the long-run. So it has a strict incentive to raise its price today in order to reduce information flows to Firm i and, hence, delay second-mover adoption. If $q(\alpha_t) < \delta(\alpha_t;\tau_i)$, then Firm j expects its quality advantage to increase in the long-run. So it has a strict incentive to lower its price today in order to increase information flows to Firm i and, hence, accelerate second-mover adoption.

The comparison of $p_i^F(\alpha_t)$ to $\bar{p}_i^F(\alpha_t)$ is less obvious because it involves both static and dynamic responses to Firm j's price. Firm i's static best-response to $p_j^L(\alpha_t)$, denoted $\overline{BR}_i^F(\alpha_t)$, is given by

$$2\overline{BR}_i^F(\alpha_t) = h + \Delta_i - q(\alpha_t) + p_j^L(\alpha_t) \iff \overline{BR}_i^F(\alpha_t) = \bar{p}_i^F(\alpha_t) - \frac{1}{3}\left(\delta(\alpha_t;\tau_i) - q(\alpha_t)\right). \quad (17)$$

Thus, Firm i has a static incentive to adjust $p_i^F(\alpha_t)$ in the same direction as $p_j^L(\alpha_t)$. Intuitively, this is because prices are strategic complements.

The dynamic portion of Firm i's best-response to $p_j^L(\alpha_t)$ is captured by the difference between $p_i^F(\alpha_t)$ and $\overline{BR}_i^F(\alpha_t)$. Upon simplification, this difference becomes

$$p_i^F(\alpha_t) - \overline{BR}_i^F(\alpha_t) = 2h\left[\frac{h + \frac{1}{3}(\Delta_i - \delta(\alpha_t;\tau_i))}{2h} - n_i^L(\alpha_t)\right]. \quad (18)$$

The bracketed expression in (18) is the difference between (i) Firm i's expected long-run market share, and (ii) Firm i's current market share. If Firm i expects its market share (and, hence, profit) to be higher in the long-run than it is today, then Firm i wants to price above $\overline{BR}_i^F(\alpha_t)$ in order to increase the Firm j's market share and accelerate learning. However, if Firm j expects its market share to be lower in the the long-run than it is today, then Firm i wants to price below $\overline{BR}_i^F(\alpha_t)$ in order to reduce Firm j's market share and decelerate learning.

2.4.2 The Second-Mover's Problem

Now take prices as given and consider Firm i's problem as a second-mover. Firm i's objective is to choose the stopping time τ_i to maximize $V_i^F(\alpha_t)$ for all $\alpha_t \in (0,1)$. To do this Firm i compares the discounted value of stopping, namely $\Pi(\Delta_i)$, to the discounted value of waiting, namely

$$\rho V_i^F(\alpha_t) = \Pi_i^F(\alpha_t) + \lambda(\alpha_t)\left[\Omega_i^F(\alpha_t) - V_i^F(\alpha_t)\right] + (\lambda_B - \lambda_G)n_j^L(\alpha_t)\alpha_t(1-\alpha_t)(V_i^F)'(\alpha_t). \tag{19}$$

Because $\Omega_i(\cdot)$ is strictly decreasing and satisfies $\rho\Omega_i^F(1) < \Pi(\Delta_i)$, it is impossible for waiting to be optimal (i.e. $\rho V_i^F(\alpha_t) > \Pi(\Delta_i)$) for all $\alpha_t \in (0,1)$ whenever Firm i is relatively patient. Instead, optimal stopping occurs when beliefs are sufficiently optimistic. Thus $\tau_i = \inf\{t \geq 0 : \alpha_t \geq \alpha_i^F\}$ for some belief $\alpha_i^F \in (0,1)$; I refer to α_i^F as Firm i's *second-mover threshold.*

Conditional on no news, posterior beliefs are increasing over time. Thus, Firm i's second-mover threshold constitutes a regular stopping threshold and is characterized by two boundary conditions: (i) $\rho V_i^F(\alpha_i^F) = \Pi(\Delta_i)$ (value matching), and (ii) $(V_i^F)'(\alpha_i^F) = 0$ (smooth pasting). Inserting these boundary conditions into (19) yields the following solution as $\rho \to 0$.

Proposition 2 (Equilibrium Stopping)**.** *Under zero discounting, Firm* i *adopts the innovation as a second-mover at time* $\tau_i = \inf\{t \geq 0 : \alpha_t \geq \alpha_i^F\}$*, where*

$$\alpha_i^F = \left[1 + \frac{\lambda_G}{\lambda_B} \frac{\Pi(\Delta_i) - \Pi(\Delta_i - x)}{\Pi(\Delta_i - q_L) - \Pi(\Delta_i)}\right]^{-1}. \tag{20}$$

To provide intuition for the above solution, note that value-matching and smooth-pasting jointly imply that Firm i's marginal benefit of waiting and marginal cost of waiting are equal at α_i^F. For any $\rho > 0$, the marginal benefit of waiting at α_i^F is

$$\alpha_i^F \lambda_B n_j^L(\alpha_i^F)\left[\frac{\Pi(\Delta_i - q_L)}{\rho} - \frac{\Pi(\Delta_i)}{\rho}\right],$$

while the marginal cost is

$$\Pi(\Delta_i) - \Pi_i^F(\alpha_i^F) + (1 - \alpha_i^F)\lambda_G n_j^L(\alpha_i^F)\left[\frac{\Pi(\Delta_i - x)}{\rho} - \frac{\Pi(\Delta_i)}{\rho}\right].$$

As discounting goes to zero, the term $\Pi(\Delta_i) - \Pi_i^F(\alpha_i^F)$ in the marginal cost of waiting becomes negligible compared to the gains or losses from news arrival. Second-mover incentives therefore depend only on $\lambda_B\,[\Pi(\Delta_i - q_L) - \Pi(\Delta_i)]$ and $\lambda_G\,[\Pi(\Delta_i) - \Pi(\Delta_i - x)]$

in the limit.

The third comparative static in Proposition 2 – that α_i^F is decreasing in Δ_i – is the least obvious but most important for my analysis. The driving force behind this comparative static is that an increase in Δ_i raises both the marginal benefit of waiting, as measured by $\Pi(\Delta_i - q_L) - \Pi(\Delta_i)$, as well as the marginal cost of waiting, as measured by $\Pi(\Delta_i) - \Pi(\Delta_i - x)$. Because flow profits are strictly convex, the increase in marginal cost always outpaces the increase in marginal benefit. Thus, the incumbent experiences a greater opportunity cost of waiting than the challenger and chooses a lower adoption threshold as a result.

The above comparative static, which I call the *cost-of-waiting effect*, clarifies an important distinction between exogenous and endogenous learning. In particular, if learning were completely exogenous (i.e. independent of first-mover sales), then Proposition 2(iii) would imply the following result:

Observation 1 (Exogenous Diffusion)**.** *Suppose that news arrival (and, hence, learning) is independent of first-mover sales. Then the incumbent will be faster (on average) than the challenger to adopt high-quality innovations as a second-mover, holding all else equal.*

The predictions of Observation 1 run counter to much empirical observation, which suggests that exogenous learning is not an ideal assumption to explain real-world patterns of diffusion. But if learning is endogenous, then equilibrium dynamics are enriched by the fact that initial quality advantage affects not only the second-mover's choice of α_i^F, but also the speed at which beliefs approach α_i^F. As I show in the next subsection, the presence of this additional economic force generates important differences in predicted equilibrium dynamics.

2.4.3 Equilibrium Dynamics

Using Propositions 1 and 2, we may now explore the implications of endogenous learning for the dynamics of price competition and second-mover adoption.

Price Competition To characterize the dynamics of price competition, it is helpful to first consider the basic properties of MPE prices when learning is exogenous. In this benchmark setting, the second-mover adoption threshold continues to satisfy Equation (20), but prices and market shares coincide with the static equilibrium:

- First-mover: $\bar{p}_j^L(\alpha) = h + \frac{1}{3}(\Delta_j + q(\alpha))$ and $\bar{n}_j^L(\alpha) = \dfrac{h + \frac{1}{3}(\Delta_j + q(\alpha))}{2h}$.
- Second-mover: $\bar{p}_i^F(\alpha) = h + \frac{1}{3}(\Delta_i - q(\alpha))$ and $\bar{n}_i^F(\alpha) = \dfrac{h + \frac{1}{3}(\Delta_i - q(\alpha))}{2h}$.

When learning is exogenous, firms can do no better than to maximize flow profit. The dynamics of price competition are therefore determined solely by the value of current beliefs. In particular, both $\bar{p}_j^L(\alpha_t)$ and $\bar{n}_j^L(\alpha_t)$ linear in α_t and increasing over time, while $\bar{p}_i^F(\alpha_t)$ and $\bar{n}_i^F(\alpha_t)$ are linear in α_t and decreasing over time. (Note: linearity is a consequence of my linear demand specification.)

When learning is endogenous, firms must account for the dynamic effect that first-mover sales have on generating information about the innovation's quality. As my next proposition shows, this additional consideration significantly alters the dynamics of competition in Stage 2.

Proposition 3 (Competition Dynamics)**.**

1. *$p_j^L(\alpha_t) - \bar{p}_j^L(\alpha_t)$ and $p_i^F(\alpha_t) - \bar{p}_i^F(\alpha_t)$ are positive over $(0, \alpha_i^F)$.*
2. *$p_j^L(\alpha_t) - \bar{p}_j^L(\alpha_t)$ is increasing and strictly convex over $(0, \alpha_i^F)$. Thus, $p_j^L(\alpha_t) - \bar{p}_j^L(\alpha_t)$ conditional on no news is increasing over time.*
3. *$n_j^L(\alpha_t) - \bar{n}_j^L(\alpha_t)$ is strictly concave in $(0, \alpha_i^F)$ and crosses zero from above. Thus, if α_0 is sufficiently close to zero, then $n_j^L(\alpha_t) - \bar{n}_j^L(\alpha_t)$ conditional on no news is (i) initially positive and increasing over time, but (ii) eventually negative and decreasing over time.*

Part (i) of Proposition 3 shows that firms possess joint incentives to learn about the innovation's quality, which leads to a softening of price competition. As in Bergemann and Välimäki (1997), this result is due the value of information being positive for each firm. The difference however is that both firms set prices that exceed the static optimum in my framework, whereas only the firm whose quality is known does in theirs. The reason for this difference is the possibility of second-mover adoption: after the arrival of good news, second-mover adoption reduces the first-mover's profit (because $x < q_H$); thus, it has less incentive to set a low price compared to when second-mover adoption is impossible.[3]

Parts (ii) and (iii) describe the implications of my framework for the dynamics of first-mover prices and market shares in Stage 2. In particular, they establish that the first-mover sets a low introductory price that sharply increases over time as the date of second-mover adoption approaches. Intuition for this penetration pricing dynamic can be developed as

---

[3] If second-mover adoption were impossible (or if $x \geq q_H$), then the prices and market shares given in Proposition 1 would reduce to exactly those in Bergemann and Välimäki (1997), and therefore the dynamics of competition over time would be completely analogous.

follows. Immediately after first-mover adoption is when beliefs about the innovation true quality are most pessimistic (conditional on no news). To attract consumers and to improve customer beliefs, the first-mover therefore sets a low initial price. At the same time, the second-mover has little incentive to prevent the first-mover from generating this information, because news is most likely to be bad news. The second-mover therefore raises its price so as to give the first-mover a high initial market share – see Figure 1.

If news does not arrive, then consumers become more optimistic about the innovation's true quality and so does the second-mover. At this point, the second-mover is unwilling to compete as softly against the first-mover because (i) the first-mover receives now a larger market share than it did before, and (ii) any news generated at this time is likely to be good news. Likewise, the first-mover is less willing to generate large amounts of information because doing so only accelerates second-mover adoption, which is unprofitable whenever $q(\alpha_t) > x$. To defend its monopoly over the innovation, the first-mover sharply raises its price so as to reduce the flow of information to the second-mover. Consequently, the first-mover's market share begins to *decrease* over time, and eventually it falls *below* the static equilibrium level as $\tau \to \tau_i$.

Innovation Diffusion To characterize the impact of endogenous learning on the dynamics of second-mover adoption, recall that $\tau_i = \inf\{t \geq 0 : \alpha_t \geq \alpha_i^F\}$ depends on Δ_i through two distinct channels. The first channel, the *cost-of-waiting effect*, which measures the direct effect of Δ_i on Firm is choice of α_i^F. The second channel, which I call the *information effect*, measures the indirect effect of Δ_i on $n_j^L(\alpha_t)$ and, hence, the amount of information

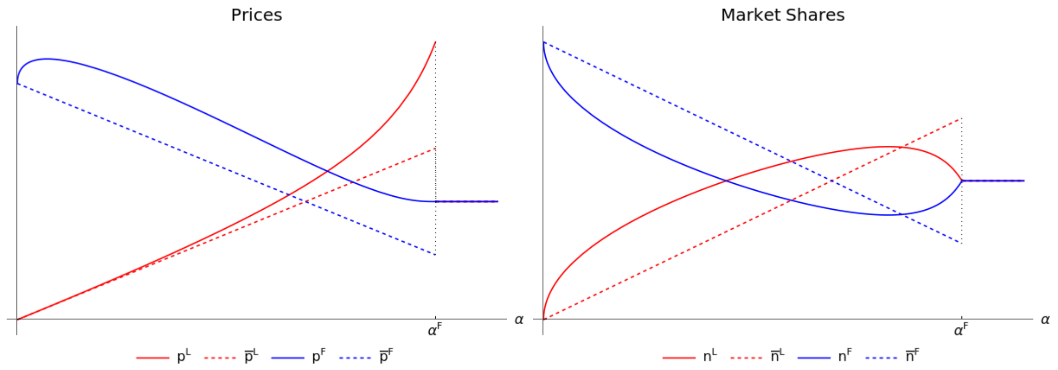

Figure 1: Equilibrium Prices and Market Shares ($q_I = q_C$)

generated about innovation quality.

All else equal, the cost-of-waiting effect is positive and therefore suggests that higher-quality firms should be faster second-movers (on average) than a lower-quality firms because they choose lower second-mover thresholds. The information effect, by contrast, is negative (see Lemma 6 in the Appendix) and therefore suggests the opposite: all else equal, higher-quality firms observe a lower volume of first-mover sales over time; therefore, it takes longer to generate good news and for posterior beliefs to reach α_i^F.

To decompose second-mover adoption incentives into the cost-of-waiting effect and information effect, respectively, let T_i denote the realized date at which Firm i adopts the innovation as a second-mover. Conditional on no news, we can rewrite T_i using the law of motion 3 and a change of variables:

$$T_i = \int_0^{T_i} dt = \int_{\alpha_0}^{\alpha_i^F} \frac{d\alpha}{(\lambda_B - \lambda_G) n_j^L(\alpha)\alpha(1-\alpha)}. \tag{21}$$

Totally differentiate this expression with respect to Δ_i uncover the two effects:

$$\frac{dT_i}{d\Delta_i} = \underbrace{\left[\frac{\partial \alpha_i^F / \partial \Delta_i}{(\lambda_B - \lambda_G) n_j^L(\alpha_i^F) \alpha_i^F (1-\alpha_i^F)}\right]}_{\text{Cost-of-Waiting Effect } (<0)} + \underbrace{\int_{\alpha_0}^{\alpha_i^F} \left[\frac{-\partial n_j^L(\alpha) / \partial \Delta_i}{(\lambda_B - \lambda_G) n_j^L(\alpha)^2 \alpha (1-\alpha)}\right] d\alpha}_{\text{Information Effect } (>0)}. \tag{22}$$

The key technical challenge in determining the sign of $dT_i/d\Delta_i$ is to compare the cost-of-waiting effect and the information effect. My first main result, Theorem 1 below, provides two sufficient conditions under which an analytical comparison can be made. In each case, the information effect turns out to be the *dominant* force in second-mover incentives. We thus obtain a complete reversal of Observation 1.

Theorem 1 (Endogenous Diffusion)**.** *Suppose that either* α_0 *is sufficiently small or that* λ_B/λ_G *is sufficiently large. Then the incumbent will be slower (on average) than the challenger to adopt high-quality innovations as a second-mover, holding all else equal.*

When learning is endogenous, second-mover adoption incentives become inextricably linked to market structure through the information effect. Theorem 1 thus provides an answer to the following historically-elusive question: *why is it so common that historically successful (i.e. high-quality) incumbents delay innovation compared to unestablished challengers* (Christensen, 1997; Cooper and Schendel, 1976; Henderson and Clark, 1990). Compared to existing theories, my framework explains this phenomenon (often called the "Innovator's Dilemma") without positing any differences between the incumbent and challenger in terms of risk aversion, organizational flexibility, or competency with new technologies. Instead, the only difference between firms is their initial quality advantage.

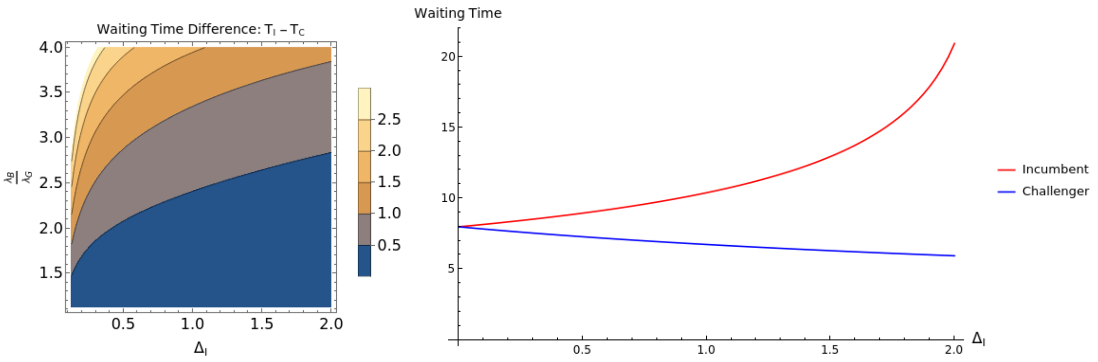

Figure 2: Second-Mover Waiting Times (No News Arrival)

For the sake of tractability, Theorem 1 is stated in terms of sufficient conditions under which the information effect and cost-of-waiting effect can be compared analytically. Numerical simulations, however, confirm that the result holds quite generally – see Figure 2. Each condition in Theorem 1 guarantees that both firms wait a sufficiently long time as second-movers; hence, the information effect is large. But importantly, the conditions I provide match the features of an empirical context in which incumbents are especially known to struggle, namely the adoption of so-called disruptive innovations (Christensen, 1993; Igami, 2017), which are typically characterized by poor initial performance (i.e. low initial beliefs) followed by rapid improvements in quality (i.e. the arrival of good news).[4]

2.4.4 The First-Mover's Problem

Turning now to initial adoption, suppose that Firm j is expected to adopt in Stage 1. Then Firm i has a strict incentive to wait (i.e. become a second-mover) if and only if the difference $F_i(\alpha_0) \equiv v_i^F(\alpha_0) - \Pi(\Delta_i)$ is positive; I refer to this as Firm i's *follower curve*.

---

[4] See Christensen (1997) for an overview of the empirical patterns associated with disruptive innovation.

Propositions 2 and 3 imply that $F_i(\alpha_0)$ is decreasing and strictly convex over $(0, \alpha_i^F)$ and satisfies $F_i(\alpha_0) > 0$ if and only if $\alpha_0 < \alpha_i^F$. Thus, Firm i's incentive to wait in Stage 1 (given adoption by Firm j) is completely summarized by the difference between α_0 and α_i^F.

Now suppose that Firm j is expected to wait in Stage 1. Then Firm i has a strict incentive to adopt (i.e. become the first-mover) if and only if the difference $L_i(\alpha_0) \equiv v_i^L(\alpha) - \Pi(\Delta_i)$ is positive; I refer to this as Firm i's *leader curve*. Propositions 2 and 3 imply that $L_i(\alpha_0)$ is strictly concave over $(0, \alpha_j^F)$ and satisfies $L_i(\alpha_0) > 0$ whenever $\alpha_0 \in (0, \alpha_j^F)$ is not too close to zero. Thus, first-mover incentives in Stage 1 possess a similar threshold characterization to second-mover incentives; namely, there exists a belief α_i^L such that $L_i(\alpha_0) > 0$ if and only if $\alpha_0 \in (\alpha_i^L, \alpha_j^F)$.

First-mover incentives in Stage 1 are jointly determined by market structure and anticipated learning effects. To see these effects, totally differentiate the indifference condition $L_i(\alpha_i^L) = 0$ with respect to Δ_i and solve for $d\alpha_i^L / d\Delta_i$:

$$\begin{aligned} L'(\alpha_i^L) \cdot \frac{d\alpha_i^L}{d\Delta_i} + \frac{\partial L_i(\alpha_i^L)}{\partial \Delta_i} + \left(\frac{\partial L_i(\alpha_i^L)}{\partial \alpha_j^F} \right) \cdot \frac{d\alpha_j^F}{d\Delta_i} &= 0 \\ \Longleftrightarrow \frac{d\alpha_i^L}{d\Delta_i} = \underbrace{\frac{\partial L_i(\alpha_i^L)/\partial \Delta_i}{-L'(\alpha_i^L)}}_{\text{Own-profit effect } (< 0)} &+ \underbrace{\frac{\partial L_i(\alpha_i^L)/\partial \alpha_j^F}{-L_i(\alpha_i^L)} \cdot \frac{d\alpha_j^F}{d\Delta_i}}_{\text{Rival-waiting effect } (> 0)} . \end{aligned} \tag{23}$$

The first term in (23) captures the direct effect of Δ_i on Firm i's relative gains/losses from initial adoption; I refer to this term as the *own-profit effect*. Because $\Pi(\cdot)$ is strictly convex, an increase in Δ_i raises both the losses from generating bad news as well as the gains from generating good news; but similar to the cost-of-waiting effect, the former dominates. The own-profit effect therefore is negative and causes a higher-quality firm to choose a *higher* first-mover threshold.

The second term in (23) captures the indirect effect of Δ_i on Firm j's second-mover incentives; I refer to this term as the *rival waiting effect*. In contrast to the own-profit effect, the rival waiting effect is *positive* and *increases* a high-quality firm's incentive to be a first-mover. The intuition behind this comparative static is the following. As a first-mover, a high-quality firm expects that its rival (due to the cost-of-waiting effect) will choose a high second-mover threshold. Both good news and bad news are thus more likely to arrive in Stage 2, which generates a mean-preserving spread over possible long-run quality advantages. Because $\Pi(\cdot)$ is strictly convex, this increase in variance is desirable and raises a high-quality firm's incentive to be a first-mover.

In general, the comparative statics of first-mover incentives may be dominated by either the own-profit effect or the rival-waiting effect – see Figure 3. This makes it difficult to analytically compare first-mover incentives of the incumbent and challenger because a local inspection of $\partial \alpha_i^L / \partial \Delta_i$ is not sufficient to infer a global comparison. My next result, Theorem 2 below, provides a sufficient condition under which an analytical comparison can be made.

Theorem 2 (Initial Adoption)**.** *Suppose that* λ_B/λ_G *is sufficiently large. Then exists a unique configuration of first-mover and second-mover thresholds, namely* $0 < \alpha_C^L < \alpha_I^L < \alpha_I^F < \alpha_C^F < 1$. *Thus, with the exception of boundary cases, there is a unique MPE outcome* *for all* $\alpha_0 \in [0, \alpha_I^L] \cup [\alpha_I^F, 1]$*:*

1. *If* $\alpha_0 \in (0, \alpha_C^L)$*, then neither firm adopts the innovation.*
2. *If* $\alpha_0 \in (\alpha_C^L, \alpha_I^L)$*, then the challenger adopts as the unique first-mover.*
3. *If* $\alpha_0 \in (\alpha_I^F, \alpha_C^F)$*, then the incumbent adopts as the unique first-mover.*
4. *If* $\alpha_0 \in (\alpha_C^F, 1)$*, then simultaneous adoption occurs.*

For $\alpha_0 \in [\alpha_I^L, \alpha_I^F]$*, there are two asymmetric pure-strategy MPE (i.e. one firm adopts, the other waits) and one mixed-strategy MPE in which Firm i adopts at rate* $\phi_i^* = L_j(\alpha_0)/(L_j(\alpha_0) + F_j(\alpha_0))$ *until one firm adopts.*

Theorem 2 obtains a unique comparison of α_I^L and α_C^L by showing that the own-profit effect always dominates the rival-waiting effect whenever λ_G/λ_B is sufficiently large. Numerical simulations, however, confirm that the result is fully general – see Figure 3. We thus obtain an equilibrium in which firms *endogenously* select into the early adoption of different classes of innovation. For relatively risky innovations (i.e. low α_0), both the incumbent and challenger are willing to wait, but only the challenger is willing to lead (due to the cost-of-waiting effect). So the unique MPE outcome is challenger-led adoption. But for relatively safe innovations (i.e. high α_0), both the incumbent and challenger are willing to wait, but only the challenger is willing to lead (due to the own-profit effect). So unique MPE outcome is incumbent-led adoption.

Unlike Theorem 1, the conclusions of Theorem 2 do not intrinsically depend on the speed of learning. This is due to my focus on zero discounting: whenever firms are infinitely patient, the expected payoff to being a first-mover depends only *how much* learning

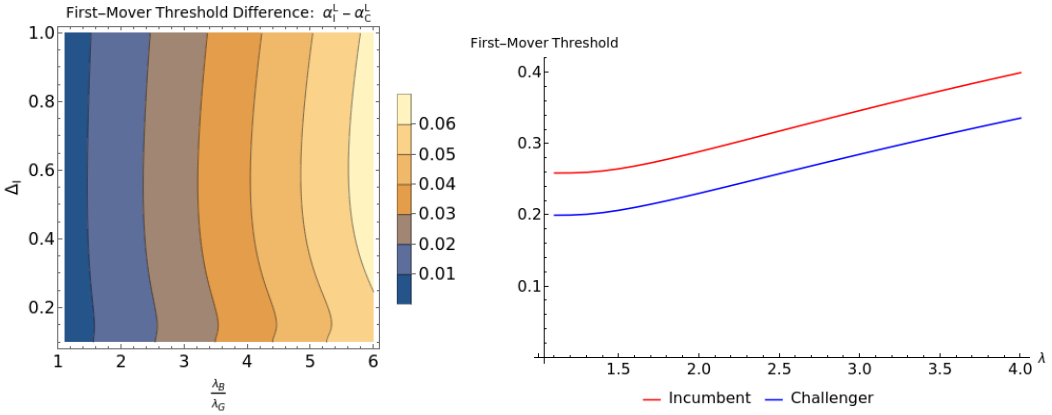

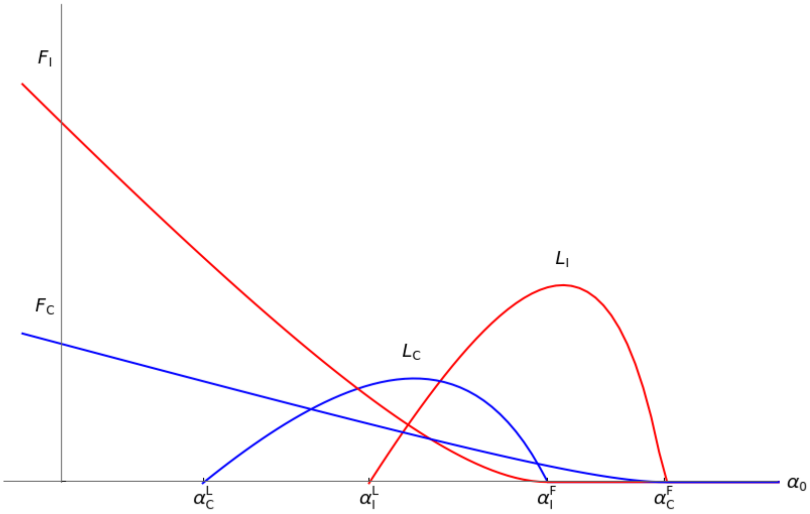

Figure 3: Comparison of Leader and Follower Curves

occurs in equilibrium, not on *how quickly* it occurs. Regardless, the initial adoption patterns described in Theorem 2 have implications for a literature that has largely focused on a dichotomous question: *are large, established incumbents more or less innovative than small, unestablished challengers?*[5] General conclusions in this area have been historically elusive, and results appear to be sensitive to industry characteristics (Gilbert, 2006c; Motta, 2004a). My results offer a possible explanation: it is not that incumbents are are more or less innovative than challengers; instead, they prioritize the development of qualitatively

---

[5] See Cohen (2010a) for a detailed survey of the empirical side of this literature.

different technologies. Incumbents are more constrained by the own-profit effect and thus prioritize safe (i.e. highly-promising) innovations, while challengers are less constrained by the own-profit effect and thus prioritize riskier (i.e. less-promising) innovations.

By combining the results of Theorem 1 and Theorem 2, we expect to observe in equilibrium that incumbents (i) tend to wait longer as second-movers, and (ii) tend to be more successful as first-movers. Elements of this first prediction have been documented in managerial studies of market leadership and innovation, particularly the vulnerability of incumbents to disruptive innovation (e.g. Christensen, 1993; Cooper and Schendel, 1976; Henderson and Clark, 1990). Elements of the second prediction have been documented in studies of increasing dominance in highly-innovative industries (e.g. Athey and Schmutzler, 2001; Bagwell and Ramey, 1994; Cabral, 2018; Klepper, 1996; Riordan and Salant, 1994). Much of this literature argues that increasing dominance can arise on the intensive margin (e.g. R&D intensity). My results offer a complementary perspective by showing that increasing dominance can also arise on the extensive margin, namely in the decision of which innovations to initially adopt.

2.5 Market Structure and Welfare

To consider the normative implications of my framework, I now analyze the potential for government regulation and changes in market structure to improve social welfare. Specifically, I compare MPE adoption and learning to that which would occur in two counterfactual scenarios: (i) if a benevolent social planner were able to directly regulate firms' pricing and adoption decisions, and (ii) if both firms were owned by a single joint-profit maximizing monopolist.

In the analysis of each counterfactual scenario, I maintain the following assumption.

Assumption 3 (Contestability). $\left|q_i - q_j\right| < \max\{q_H, -q_L\}$.

Assumption 3 strengthens Assumption 1 to require that (i) the challenger can surpass the incumbent in terms of quality as the sole-adopter a high-quality innovation, and (ii) the incumbent can fall behind the challenger in terms of quality as the sole-adopter of a low-quality innovation. In other words, $q_I - q_C$ is sufficiently close to zero that reversals in quality advantage can occur. This simplifies my exposition by reducing the number of possible cases to consider.

2.5.1 The Planner's Solution

Begin by considering the problem of a social planner with the power to regulate firms' adoption decisions and market shares via taxes and/or subsidies.

Second-Mover Adoption. If Firm j has adopted the innovation but Firm i has not, then the planner will choose a belief threshold $\alpha^F_{i,S}$ above which Firm i should adopt as a second-mover. Under zero discounting, the solution to the planner's stopping problem will be expressed in terms of static flow total surplus. So it is helpful to first record the following lemma.

Lemma 2. *The statically-optimal market share for Firm i at time t is* $n_{i,S}(\Delta_{i,t}) = (h +$

$\Delta_{i,t})/(2h)$. *The resulting flow total surplus is*

$$TS(q_{i,t},q_{j,t}) = q_{i,t}n_S(\Delta_{i,t}) + q_{j,t}n_S(\Delta_{j,t}) + h\left[\frac{1}{2} + n_S(\Delta_{i,t})\big(1 - n_S(\Delta_{i,t})\big)\right].$$

The planner's objective is to maximize long-run average total surplus, which can be expressed for a particular value of $\alpha^F_{i,S}$ as follows:

$$\begin{aligned} v^F_{i,S}(\alpha_t) \equiv TS(q_i,q_j) + q(\alpha_t) \\ + \theta_G(\alpha_t;\alpha^F_{i,S})\big[TS(q_i,q_j - x) - TS(q_i,q_j)\big] \\ + \theta_B(\alpha_t;\alpha^F_{i,S})\big[TS(q_i + q_L,q_j) - TS(q_i,q_j) - q_L\big], \end{aligned} \tag{24}$$

where $\theta_G(\alpha_t;\alpha^F_{i,S})$, $\theta_B(\alpha_t;\alpha^F_{i,S})$ respectively denote the probability that good news, bad news arrives before second-mover adoption occurs. The first-order condition defining $\alpha^F_{i,S}$ is analogous to that which defines α^F_i. At the margin, the social benefit of additional waiting, namely

$$\frac{\partial\theta_B(\alpha_t;\alpha^F_{i,S})}{\partial\alpha^F_{j,S}}\big[TS(q_i,q_j + q_L) - TS(q_i,q_j) - q_L\big], \tag{25}$$

must coincide with the social cost of additional waiting, namely

$$\frac{\partial\theta_G(\alpha_t;\alpha^F_{i,S})}{\partial\alpha^F_{j,S}}\big[TS(q_i,q_j) - TS(q_i - x,q_j)\big]. \tag{26}$$

Equating (25) and (26), and then solving for $\alpha^F_{i,S}$ yields the following analytical solution.

Proposition 4 (Optimal Optimal Stopping)**.** *Suppose that Firm i is the second-mover in*

Stage 2. Then the planner's optimal choice of second-mover threshold is

$$\alpha^F_{i,S} = \left[1 + \frac{\lambda_G}{\lambda_B} \frac{TS(q_i, q_j) - TS(q_i - x, q_j)}{TS(q_i, q_j + q_L) - TS(q_i, q_j) - q_L} \right]^{-1}. \tag{27}$$

Compared to the MPE, we have $\alpha^F_{i,S} > \alpha^F_i$. Thus, the planner has a strict incentive to subsidize addition waiting by second-movers in any MPE.

The solution for $\alpha^F_{i,S}$ has a similar interpretation to α^F_i. The clear difference of course is that $\alpha^F_{i,S}$ is determined by the social costs and benefit of additional waiting. Any disagreement between MPE and socially-optimal waiting can therefore be attributed to the fact that the planner (i) internalizes the effects of technology adoption on consumer surplus and (ii) does not care about the relative distribution of first-mover and second-mover profits. All else equal, the first consideration raises the planner's incentive to wait for bad news, while the second raises the planner's incentive to wait for good news. The overall effect is that the planner is unambiguously more willing than Firm i to wait for news in equilibrium.

Speed of Learning. Having determined the planner's choice of $\alpha^F_{i,S}$, we can now determine the optimal speed of learning in Stage 2. Once again, I assume the planner evaluates total surplus using the strong long-run average criterion so that we can obtain an analytically tractable version of the planner's experimentation program, namely

$$n^L_{j,S}(\alpha_t) = \arg\max_{n_i} \left\{ \frac{(q_j + q(\alpha_t))n_j + q_i(1 - n_j) + \frac{h}{2} + hn_j(1 - n_j) - v^L_{j,S}(\alpha_t)}{n_j} \right\}. \tag{28}$$

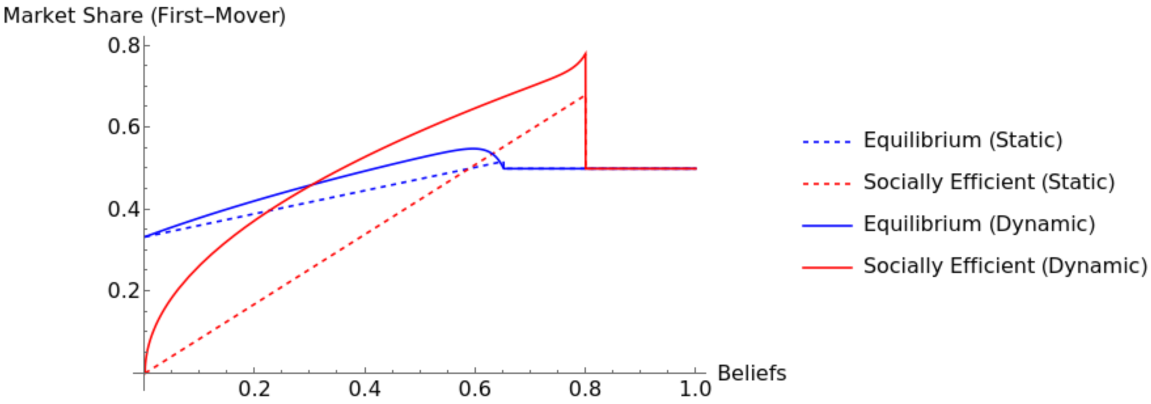

Figure 4: Equilibrium vs Optimal Experimentation

Note that (15) is directly analogous to (12) and (12); in particular, the planner maximizes lifetime total surplus net long-run average total surplus. The optimal choice of $n^L_{j,S}(\alpha_t)$ therefore balances static and dynamic considerations in a similar manner to the positive-discounting case.

Proposition 5 (Optimal Experimentation)**.** *Suppose that Firm j is the first-mover in the planner's optimal adoption policy. Then the socially optimal intensity of experimentation in Stage 2 is*

$$n^L_{i,S}(\alpha_t) = \sqrt{\frac{\nu^L_{j,S}(\alpha_t) - q_i - h/2}{h}}. \tag{29}$$

Whenever λ_B/λ_G is sufficiently large, the difference $n^L_{i,S}(\alpha_t) - n^L_i(\alpha_t)$ is increasing over $(0, \alpha^F_j)$ and crosses zero once.

Figure 4 depicts the optimal experimentation policy compared to MPE experimentation. As in the case of MPE experimentation, the socially-optimal policy converges to the static optimum as $\alpha_t \to 0$ because the incentive to experiment vanishes when uncertainty is resolved. But the same cannot be said as $\alpha_t \to \alpha^L_{i,S}$. In particular, the planner's optimal choice of $n^L_{i,S}(\alpha_t)$ is discontinuous at $\alpha^L_{i,S}$ because there is positive social value of information even as the date of second-mover adoption approaches. This feature of optimal experimentation stands in sharp contrast to the case of MPE experimentation, where all first-mover incentives to experiment vanish as $t \to \tau_i$. In fact, numerical simulations indicate that the social benefit of experimentation *increases* as the date of second-mover adoption because the planner has large incentives to avoid having the second-mover mistakenly a low-quality innovation.

In general, the planner would like first-movers to experiment more intensely as beliefs become more optimistic. Therefore, the penetration pricing dynamics identified in Proposition 3 are socially inefficient. In particular, MPE experimentation is always *higher* than the social optimum when α_t is relatively small and *lower* than the social optimum when α_t is relatively large. This reversal in the direction of social inefficiency occurs because the first-mover wishes to generate (i) large amounts of information early in Stage 2, but (ii) small amounts of information as the date of second-mover adoption approaches. First-movers therefore have inefficiently high incentives to experiment when α_t is low and inefficiently low incentives to experiment when α_t is high.

Initial Adoption. To characterize socially-efficient initial adoption, define Firm i's *social leader curve* as $L_{i,S}(\alpha_0) \equiv v^L_{i,S}(\alpha_0) - TS(q_i, q_j)$ and its *social follower curve* as $F_{i,S}(\alpha_0) = v^F_{i,S}(\alpha_0) - TS(q_i, q_j) - q(\alpha_0)$. By symmetry, social leader and follower curves always satisfy the identity $L_{i,S}(\alpha_0) = F_{j,S}(\alpha_0) + q(\alpha_0)$. We therefore need only consider social leader curves.

There are two possible sources of inefficiency in equilibrium initial adoption. First, the order of adoption is inefficient whenever Firm i adopts the innovation as a first-mover but Firm j would have been a more efficient first-mover; that is, $L_{i,S}(\alpha_0) < L_{j,S}(\alpha_0)$. Second, the timing of adoption is inefficient whenever Firm i adopts as a first-mover and we have either $L_{i,S}(\alpha_0) < 0$ (i.e. non-adoption is optimal) or $F_{j,S}(\alpha_0) < 0$ (i.e. simultaneous adoption is optimal). The following lemma characterizes the optimal order of initial adoption.

Lemma 3. *Suppose that λ_B/λ_G is sufficiently large. Then there exists a unique belief $\hat{\alpha}_S \in (0, \alpha^F_{I,S})$ such that the incumbent is the socially-preferred first-mover if and only if $\alpha_0 > \hat{\alpha}_S$.*

An interesting feature of efficient sequential adoption is that the choice of first-mover is (i) always unique, and (ii) changes depending on whether α_0 is above or below a pivotal belief $\hat{\alpha}_S$. Intuitively, the planner faces a dynamic trade-off when deciding the order of adoption. On the one hand, initial adoption by the incumbent ensures that a large number of consumers receive the innovation's expected benefit in Stage 2, and that learning occurs rather quickly. On the other hand, such adoption is risky as the innovation may be low-quality and generate large welfare losses. The planner resolves this trade-off by having the high-quality firm lead adoption if and only if α_0 is sufficiently high – namely, above $\hat{\alpha}_S$.

Having determined the optimal order of initial adoption, we may now precisely characterize the planner's optimal choice of timing.

Proposition 6 (Optimal Initial Adoption)**.** *Suppose that λ_B/λ_G is sufficiently large. Then, with the exception of boundary cases, the social planner's optimal adoption policy in Stage*

1 is uniquely characterized by a tuple of beliefs, $(\alpha_S^L, \hat{\alpha}_S, \hat{\alpha}_S^L)$, such that:

1. *If $\alpha_0 < \alpha_S^L$, then neither firm adopts the innovation.*
2. *If $\alpha_0 \in (\alpha_S^L, \hat{\alpha}_S)$, then the challenger adopts as the unique first-mover.*
3. *If $\alpha_0 \in (\hat{\alpha}_S, \alpha_S^F)$, then the incumbent adopts as the unique first-mover.*
4. *If $\alpha_0 > \alpha_S^F$, then simultaneous adoption occurs.*

Compared to MPE initial adoption, we have $\alpha_S^L < \min\{\alpha_i^L, \alpha_j^L\}$. Thus, the planner has a strictly positive incentive to subsidize challenger-led adoption.

As in the case of MPE adoption, the socially-optimal policy favors sequential adoption over a range of intermediate beliefs. However, the planner strictly prefers that adoption occur at lower beliefs than occurs in MPE. The intuition behind this disagreement is that the planner fully internalizes the value of information from first-mover sales. In other words, firms have socially *insufficient* first-mover incentives because they undervalue the information that sequential adoption generates for the second-mover.

2.5.2 The Monopoly Solution

To analyze the impact of market structure on adoption and learning incentives, suppose that I and C merge to become a joint-profit-maximizing monopolist. To obtain the most direct comparison between the monopoly, duopoly, and socially-optimal regimes, suppose also that (i) the entire market continues to be served, and (ii) both products continue to be supplied in positive quantity. Assumption 1 is sufficient to ensure that (i) holds. To ensure that (ii) holds, it suffices to assume that the challenger's post-adoption quality always remains positive.

Assumption 4. $q_C > -q_L$.

Lemma 4. *The monopolist's statically-optimal market share and price for Firm* i *is given by* $n_{i,M}(\Delta_{i,t}) = (2h+\Delta_{i,t})/(4h)$ *and* $p_{i,M}(q_{i,t}, \Delta_{i,t}) = q_{i,t} + h(1-n_{i,M}(\Delta_{i,t}))$. *The resulting flow profit is*

$$\Pi_M(q_{i,t}, q_{j,t}) = \frac{4h(h+q_{i,t}+q_{j,t}) + \Delta_{i,t}^2}{8h}.$$

We once again proceed via backward induction, starting with the second-mover's problem.

Proposition 7 (Monopoly Stopping)**.** *Suppose that Firm* i *is the second-mover in Stage 2. Then monopolist's optimal choice of second-mover threshold is*

$$\alpha_{i,M}^F = \left[1 + \frac{\lambda_G}{\lambda_B} \frac{\Pi_M(q_i, q_j) - \Pi_M(q_i - x, q_j)}{\Pi_M(q_i, q_j + q_L) - \Pi_M(q_i, q_j) - q_L}\right]^{-1}.$$

Compared to the MPE and social optimum, we have $\alpha_i^F < \alpha_{i,M}^F < \alpha_{i,S}^F$. *Thus, monopoly waiting incentives are a strict improvement over MPE waiting.*

Proposition 7 establishes that second-mover incentives under monopoly are unambiguously more efficient than duopoly incentives. The reason is that a monopolist has the ability to coordinate prices and therefore capture a greater share of total surplus. Monopoly stopping incentives therefore internalize more of the social benefits and social costs of waiting. Nevertheless, some inefficiency remains because the monopolist does not capture the entirety of total surplus.

A general theme of the monopoly solution is that greater internalization of social surplus will produce efficiency gains whenever the direction of inefficiency under duopoly is *unambiguous*. In situations where inefficiency exists in both directions, such as the intensity of experimentation in Stage 2, monopoly incentives can be more or less efficient than duopoly incentives.

Proposition 8 (Monopoly Experimentation)**.** *Suppose that Firm j is the first-mover in Stage 2. Then the monopolist's optimal level of experimentation is*

$$n^L_{j,M}(\alpha_t) = \sqrt{\frac{v^L_{j,M}(\alpha_t) - q_i}{2h}}.$$

Whenever λ_B/λ_G *is sufficiently large, there exists a nonempty interval of beliefs* $(\underline{\alpha}, \overline{\alpha}) \subsetneq (0,1)$ *such that duopoly experimentation is closer to the social optimum than is monopoly experimentation iff* $\alpha \in (\underline{\alpha}, \overline{\alpha})$*.*

The intuition behind Proposition 8 can be understood as follows. Because the monopolist internalizes a larger fraction of the social benefits of experimentation in Stage 2, it chooses more efficient market shares over a wide range of beliefs. In particular, the monopolist prefers a lower rate of experimentation when α_t is low (efficient) and a higher rate of experimentation when α_t is high (efficient). But because the monopolist does not fully internalize the social benefits of experimentation, there always exists an intermediate range of beliefs where the the monopolist strictly prefers a higher rate of experimentation, but while the planner strictly prefers a lower rate of experimentation. In this case, monopoly

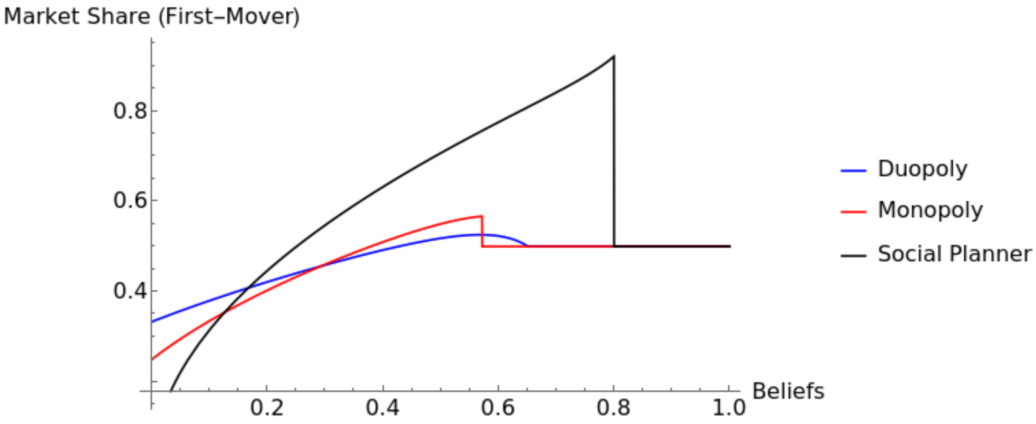

Figure 5: Monopoly, Duopoly, and Socially Efficient Experimentation

experimentation is less efficient than duopoly experimentation – see Figure 5.

Finally, initial adoption incentives in Stage 1 may be characterized via monopoly leader and follower curves just as before; denote them $L_{i,M}(\alpha_0)$ and $F_{i,M}(\alpha_0)$. As in the planner's problem, the monopolist coordinates the optimal order of sequential adoption by comparing $L_{i,M}(\alpha_0)$ to $L_{j,M}(\alpha_0)$. After determining the optimal order of adoption, the monopolist then determines whether sequential adoption is preferred to either non-adoption or simultaneous adoption.

Proposition 9 (Monopoly Initial Adoption)**.** *Suppose that λ_B/λ_G is sufficiently large. Then, with the exception of boundary cases, the monopolist's optimal adoption policy in Stage 1 is uniquely characterized by a tuple of beliefs, $(\alpha_M^L, \hat{\alpha}_M, \alpha_M^F)$, such that:*

1. *If $\alpha_0 \in (0, \alpha_M^L)$, then neither firm adopts the innovation.*
2. *If $\alpha_0 \in (\alpha_M^L, \hat{\alpha}_M)$, then the challenger adopts as the unique first-mover.*
3. *If $\alpha_0 \in (\hat{\alpha}_M, \alpha_M^F)$, then the incumbent adopts as the unique first-mover.*
4. *If $\alpha_0 > \alpha_M^F$, then simultaneous adoption occurs.*

Compared to MPE and socially-efficient initial adoption, have $\alpha_S^L < \alpha_M^L < \alpha_C^L$. *Thus, the timing of initial adoption under monopoly is a strict improvement over duopoly adoption.*

Monopolization leads to a strict improvement in the efficiency of initial adoption because the directional inefficiency of duopoly incentives is unambiguous. For this same reason, the order of initial adoption may become more or less efficient. The monopolist is *guaranteed* to choose an inefficient order of adoption for all initial beliefs satisfying $\alpha_0 \in (\min\{\hat{\alpha}_S, \hat{\alpha}_M\}, \max\{\hat{\alpha}_S, \hat{\alpha}_M\})$. Interestingly, the magnitude of this order inefficiency depends on how much learning occurs in Stage 2. In particular, it can be shown that $|\hat{\alpha}_M - \hat{\alpha}_S|$ converges to zero as $\lambda_B/\lambda_G \to \infty$. Thus, the order of monopoly adoption is most efficient whenever first-mover sales are highly informative about innovation quality.

2.6 Concluding Remarks

This dissertation chapter develops a framework to analyze the role of endogenous learning in the adoption and diffusion of an uncertain innovation. Firms learn the innovation's true quality by observing first-mover sales, which gives rise to a rich set of equilibrium dynamics. First, the information generated by first-mover sales raises consumers' expectations about innovation quality but also accelerates second-mover adoption. To balance these dual considerations, first-movers engage in *penetration pricing*: setting a low initial price that sharply increases over time so as to reduce information flows to the second-mover. Second, initial quality advantage determines the speed of innovation diffusion through both an *opportunity-cost effect* as well as an *information effect*. My framework thus provides a unified explanation for why high-quality incumbents can be highly innovative as first-movers but also disproportionately slow to adopt innovations as second-movers. Third, incumbents and challengers endogenously select into distinct innovation timing strategies based on expectations about innovation quality. Incumbents are most likely to lead in developing highly-promising innovations, while challengers are most likely to lead in developing less-promising innovations. Furthermore, the endogenous selection that occurs in equilibrium favors *increasing dominance*.

My framework is designed to be flexible and can be extended in several directions. An immediate extension we might consider is a relaxation of my assumption that no news is good news. In an alternative scenario where no news is bad news, Theorems 1 and 2 both continue to hold; in fact, they are easier to prove. The reason is that second-mover adoption takes a remarkably simple form in this case: always wait for good news. Therefore, the opportunity-cost effect vanishes from the second-mover's problem, and all that remains is the information effect. Likewise, the rival-waiting effect vanishes from the first-mover's problem, and all that remains is the own-profit effect. The main difference that emerges when no news is bad news is that Proposition 3 is effectively reversed. Beliefs are no longer increasing over time, hence first-movers no longer have strategic incentives to sharply raise prices over time. Instead, they defend their monopoly over the innovation via *price skimming*: setting a high initial price that sharply decreases over time.

Other extensions would require additional work but could yield new insights. First, by focusing on the case of a single innovation, my analysis does not consider the possibility of sequential innovation (c.f. Green and Scotchmer, 1995) or how it impacts equilibrium adoption and learning incentives. Second, the source of first-mover advantage in my framework is taken to be exogenous. We might be interested in knowing how an endogenous source of first-mover advantage, such as learning-by-doing (c.f. Cabral and Riordan, 1994), interacts with endogenous learning incentives. Such an extension would greatly complicate the analysis of the second-mover's problem and equilibrium pricing incentives, but tractability might be preserved if one considers appropriate simplifications of the base model.

2.7 Appendix: Proofs and Derivations

Proof of Lemma 1. Given $(q_{1,t}, q_{2,t})$ and $(p_{1,t}, p_{2,t})$, a consumer with type θ prefers to buy Firm 1's product if and only if $q_{1,t} - p_{1,t} + h(1-\theta) \geq q_{2,t} - p_{2,t} + h\theta \iff \theta \leq \bar{\theta} = (h + q_{1,t} - q_{2,t} + p_{2,t} - p_{1,t})/(2h)$. Provided that $\bar{\theta}$ is interior, Firm i's flow profit, namely $p_{i,t}(h + q_{i,t} - q_{j,t} + p_{j,t} - p_{i,t})/(2h)$, is maximized at $p_{i,t} = (h + q_{i,t} - q_{j,t} + p_{j,t})/2$. By symmetry, Firm j's flow profit is maximized at $p_{j,t} = (h + q_{j,t} - q_{i,t} + q_{i,t})/2$. Solving this system of best-response conditions yields $p_{i,t} = h + (q_{i,t} - q_{j,t})/3$ and $p_{j,t} = h + (q_{j,t} - q_{i,t})/3$. Thus, Firm i's market share is $n_{i,t} = (h + (q_{i,t} - q_{j,t})/3)/(2h)$ and its flow profit is $\Pi_{i,t} = (h + (q_{i,t} - q_{j,t})/3)^2/(2h)$. □

Proof of Proposition 1. The first-order condition defining Firm j's best-response as the first-mover is

$$n_j^L(\alpha)\left(n_j^L(\alpha) - p_j\left(\frac{-1}{2h}\right)\right) + \left(p_j n_j^L(\alpha) - v_j^L(\alpha)\right)\left(\frac{-1}{2h}\right) = 0,$$

which solves to give $n_j^L(\alpha) = \sqrt{\nu_j^L(\alpha)/(2h)}$. Thus, (15) holds. The first-order condition defining Firm i's best-response as the second-mover is

$$n_j^L(\alpha)\left(n_i^F(\alpha)+p_i\left(\frac{-1}{2h}\right)\right)-\left(p_i n_i^F(\alpha)-\nu_i^F(\alpha)\right)\left(\frac{1}{2h}\right)=0,$$

which solves to give $p_i = \nu_i^F(\alpha) + 2hn_j^L(\alpha)n_i^F(\alpha)$. To obtain $p_i^F(\alpha)$, insert (15) into this expression for p_i to obtain $p_i^F(\alpha) = \nu_i^F(\alpha) - \nu_j^L(\alpha) + \sqrt{2h\,\nu_j^L(\alpha)}$ which, upon simplification, yields (14). Finally, to obtain $p_j^L(\alpha)$, use the definition of the marginal consumer to obtain

$$n_j^L(\alpha)=\frac{h+\Delta_j+q(\alpha)+p_i^F(\alpha)-p_j^L(\alpha)}{2h} \iff p_j^L(\alpha)=h+\Delta_j+q(\alpha)+p_i^F(\alpha)-2hn_j^L(\alpha),$$

Insert (15) and (14) into the above expression for $p_j^L(\alpha)$, and then simplify, to obtain (13).

□

Proof of Proposition 2. Together, value-matching and smooth-pasting conditions imply

$$\begin{aligned}\Pi(\Delta_i) = \Pi_i^F(\alpha_i^F) + \alpha_i^F\lambda_G n_j^L(\alpha_i^F)\left[\frac{\Pi(\Delta_i - x)-\Pi(\Delta_i)}{\rho}\right]& \\ +(1-\alpha_i^F)\lambda_B n_j^L(\alpha_i^F)\left[\frac{\Pi(\Delta_i - q_L)-\Pi(\Delta_i)}{\rho}\right]&,\end{aligned} \tag{30}$$

which yields the following implicit expression for α_i^F:

$$\alpha_i^F = \frac{\rho\left(\frac{\Pi_i^F(\alpha_i^F)-\Pi(\Delta_i)}{n_j^L(\alpha)}\right)+\lambda_B\left[\Pi(\Delta_i - q_L)-\Pi(\Delta_i)\right]}{\lambda_B\left[\Pi(\Delta_i - q_L)-\Pi(\Delta_i)\right]+\lambda_G\left[\Pi(\Delta_i)-\Pi(\Delta_i - x)\right]}.$$

In any MPE, $n_j^L(\cdot)$ and $\Pi_i^F(\cdot)$ are uniformly-bounded. Thus, $\rho(\Pi_i^F(\alpha_i^F)-\Pi(\Delta_i))/n_j^L(\alpha_i^F) \to$

0 as $\rho \to 0$ provided that $n_j^L(\alpha_i^F) > 0$ holds in the limit, which is implied by Assumption 1. Thus, (20) holds. □

Proof of Proposition 3. Before proving Proposition 3, it is helpful to establish the following lemma.

Lemma 5. *(i)* $\theta_G(\alpha; \tau_i^F)$ *is strictly concave over* $(0, \alpha_i^F)$*, and (ii)* $\theta_B(\alpha; \tau_i)$ *is strictly convex over* $(0, \alpha_i^F)$.

Proof. The proof is by calculation. Upon simplification, the second derivative of $\theta_G(\alpha; \tau_i)$ is

$$\left(\frac{\lambda_G}{\lambda_B - \lambda_G}\right)\left(\frac{\Phi(\alpha_i^F)}{\Phi(\alpha)}\right)^{\frac{\lambda_G}{\lambda_B - \lambda_G}}\left(\frac{1}{1-\alpha}\right)\underbrace{\left[\frac{-1}{1-\alpha} - \frac{1}{\alpha}\left(1 + \left(\frac{1}{1-\alpha}\right)\left(\frac{\lambda_G}{\lambda_B - \lambda_G}\right)\right)\right]}_{\text{Negative}} < 0.$$

So $\theta_G(\alpha; \tau_i)$ is strictly concave over $(0, \alpha_i^F)$.

Likewise, the second derivative of $\theta_B(\alpha; \tau_i)$ simplifies to give

$$\theta_B''(\alpha; \tau_i) = \frac{1}{1-\alpha}\left(\frac{\lambda_B}{\lambda_B - \lambda_G}\right)\left(\frac{\Phi(\alpha_i^F)}{\Phi(\alpha)}\right)^{\frac{\lambda_B}{\lambda_B - \lambda_G}}\left(\frac{\lambda_B}{\lambda_B - \lambda_G} - 1\right)\frac{1}{(1-\alpha)\alpha} > 0.$$

So $\theta_B(\alpha; \tau_i)$ is strictly convex over $(0, \alpha_i^F)$. □

Corollary 1. *Both* $\delta(\alpha; \tau_i)$ *and* $v_j^L(\alpha)$ *are strictly concave over* $(0, \alpha_i^F)$.

Now to prove Proposition 3. To establish that $p_j^L(\alpha) - \bar{p}_j^L(\alpha)$ is positive over $(0, \alpha_i^F)$, observe that $q(\alpha) > \delta(\alpha;\tau_i)$ holds whenever

$$\begin{aligned}
\alpha x + (1-\alpha) q_L > \delta(\alpha;\tau_i) &\iff \alpha \left(\frac{\Phi(\alpha_i^F)}{\Phi(\alpha)} \right)^{\frac{\lambda_G}{\lambda_B - \lambda_G}} x > (1-\alpha) \left(\frac{\Phi(\alpha_i^F)}{\Phi(\alpha)} \right)^{\frac{\lambda_B}{\lambda_G - \lambda_B}} (-q_L) \\
&\iff \frac{x}{-q_L} > \Phi(\alpha_i^F) = \frac{\lambda_G}{\lambda_B} \frac{\Pi(\Delta_i) - \Pi(\Delta_i - x)}{\Pi(\Delta_i - q_L) - \Pi(\Delta_i)} \\
&\iff \frac{x}{-q_L} > \frac{\lambda_G}{\lambda_B} \frac{(6h + 2\Delta_i - x)x}{(6h + 2\Delta_i - q_L)(-q_L)}.
\end{aligned}$$

Since $\lambda_G < \lambda_B$, the above equality holds. To establish that $p_i^F(\alpha) - \bar{p}_i^F(\alpha)$ is positive over $(0, \alpha_i^F)$, re-write the difference as follows:

$$p_i^F(\alpha) - \bar{p}_i^F(\alpha) = \left[\sqrt{2h v_j^L(\alpha)} - h - \frac{1}{3} \left(\Delta_j + \delta(\alpha;\tau_i) \right) \right] + \frac{1}{3} \left(q(\alpha) - \delta(\alpha;\tau_i) \right).$$

The second term is positive by the previous inequality. The first term is positive because $\Pi(\cdot)$ is strictly convex:

$$\sqrt{2h v_i^L(\alpha)} > \sqrt{2h \Pi(\Delta_j + \delta(\alpha;\tau_i))} = h + \frac{1}{3} \left(\Delta_j + \delta(\alpha;\tau_i) \right).$$

This completes the proof of Part 1 of Proposition 3.

To establish Part 2 of Proposition 3, it suffices to prove that $p_i^L(\alpha) - \bar{p}_i^L(\alpha)$ is increasing over $(0, \alpha_i^F)$. To do this, first use Lemma 5 to conclude that $p_j^L(\alpha)$ is strictly convex over $(0, \alpha_i^F)$. Thus, it suffices to show that $(p_j^L)'(0) - (\bar{p}_j^L)'(0)$ is non-negative. Through direct calculation, it can be shown that

$$\begin{aligned}\delta'(\alpha;\tau_i) = &\left[1-\left(1+\left(\frac{1}{1-\alpha}\right)\left(\frac{\lambda_G}{\lambda_B-\lambda_G}\right)\right)\left(\frac{\Phi(\alpha_i^F)}{\Phi(\alpha)}\right)^{\frac{\lambda_G}{\lambda_B-\lambda_G}}\right]x \\ &+\left[1-\left(1-\left(\frac{1}{1-\alpha}\right)\left(\frac{\lambda_B}{\lambda_B-\lambda_G}\right)\right)\left(\frac{\Phi(\alpha_i^F)}{\Phi(\alpha)}\right)^{\frac{\lambda_B}{\lambda_B-\lambda_G}}\right](-q_L).\end{aligned}$$

Evaluating this derivative at $\alpha = 0$ yields $\delta'(0;\tau_i) = x - q_L = q'(0)$. Thus, $(p_j^L)'(0) - (\bar{p}_j^L)'(0) = 0$.

Finally, to establish Part 3 of Proposition 3, it suffices to prove that $n_i^L(\alpha) - \bar{n}_j^L(\alpha)$ is (i) positive and increasing whenever α is sufficiently close to zero, and (ii) negative and decreasing whenever α is sufficiently close to α_i^F. Through direct calculation, it can be shown that

$$(n_j^L)'(0) = \frac{\Pi(\Delta_j + x) - \Pi(\Delta_j + q_L)}{4hn_j^L(0)} = \frac{x-q_L}{6h}\left(\frac{6h+2\Delta_j+x+q_L}{6h+2\Delta_j+2q_L}\right) > \frac{x-q_L}{6h} = (\bar{n}_j^L)'(0).$$

Thus, $(n_j^L)'(\alpha) - (\bar{n}_j^L)'(\alpha)$ is positive whenever α is sufficiently close to zero. Furthermore, note that $\nu_j^L(\alpha) \to \Pi(\Delta_j + q_L)$ as $\alpha \to 0$. Therefore, (i) holds. To establish (ii), it suffices to prove that $n_j^L(\alpha) - \bar{n}_j^L(\alpha)$ is strictly concave over $(0, \alpha_i^F)$ and $n_j^L(\alpha_i^F -) - \bar{n}_j^L(\alpha_i^F -) < 0$. Strict concavity is implied by the strict concavity of $\nu_j^L(\alpha)$. Negativity of $n_j^L(\alpha) - \bar{n}_j^L(\alpha)$ as $\alpha \to \alpha_j^F$ is implied by continuity; specifically, $\nu_j^L(\alpha) \to \Pi(\Delta_j) < \Pi(\Delta_j + q(\alpha_i^F)$ as $\alpha \to \alpha_j^F$. □

Before proving Theorem 1, it is helpful to first establish the following market-share comparison.

Lemma 6. *$n_I^L(\alpha) > n_C^L(\alpha)$ for all $\alpha \in (0, \alpha_I^F)$.*

Proof. It suffices to show that $v_I^L(\alpha) > v_C^L(\alpha)$ holds for all $\alpha \in (0, \alpha_I^F)$. Begin by computing

$$\begin{aligned}\frac{\partial v_i^L(\alpha)}{\partial \alpha_j^F} = &\left(\frac{1}{1-\alpha}\right)\left(\frac{\lambda_G}{\lambda_B-\lambda_G}\right)\left(\frac{\Phi(\alpha_j^F)}{\Phi(\alpha)}\right)^{\frac{\lambda_G}{\lambda_B-\lambda_G}}\left[\Pi(\Delta_i+x)-\Pi(\Delta_i)\right] \\ &+\left(\frac{1}{\alpha}\right)\left(\frac{\lambda_B}{\lambda_B-\lambda_G}\right)\left(\frac{\Phi(\alpha_j^F)}{\Phi(\alpha)}\right)^{\frac{\lambda_B}{\lambda_B-\lambda_G}}\left[\Pi(\Delta_i+q_L)-\Pi(\Delta_i)\right].\end{aligned} \tag{31}$$

This derivative is positive iff

$$\left(\frac{\lambda_G}{\lambda_B}\right)\frac{\Pi(\Delta_i+x)-\Pi(\Delta_i)}{\Pi(\Delta_i)-\Pi(\Delta_i+q_L)} > \Phi(\alpha_j^F) \iff \frac{6h+2\Delta_i+x}{6h+2\Delta_i+q_L} > \frac{6h+2\Delta_j-x}{6h+2\Delta_j-q_L},$$

which holds iff $12h(x-q_L) > 0$. Thus, we conclude that $v_I^L(\alpha)$ is greater than

$$v_I^L(\alpha;\tau_I) \equiv \Pi(\Delta_I) + \theta_G(\alpha;\tau_I)\left[\Pi(\Delta_I+x)-\Pi(\Delta_I)\right] + \theta_B(\alpha;\tau_I)\left[\Pi(\Delta_I+q_L)-\Pi(\Delta_I)\right].$$

Subtracting $v_C^L(\alpha)$ from $v_I^L(\alpha;\tau_I)$ yields the following upon simplification:

$$v_I^L(\alpha;\tau_I) - v_C^L(\alpha) = \frac{2\Delta_I}{9h}\left(3h+\delta(\alpha;\tau_I)\right).$$

The above is positive under Assumption 1. Thus, $v_I^L(\alpha) > v_C^L(\alpha)$ holds for all $\alpha \in (0, \alpha_I^F)$, as desired. □

Now to prove Theorem 1.

Proof of Theorem 1. Conditional on no news, the total amount of first-mover sales required before second-mover adoption is

$$N_i^F = \frac{1}{\lambda_B - \lambda_G}\left[\log(\Phi(\alpha_0)) - \log\left(\Phi(\alpha_i^F)\right)\right].$$

As $\alpha_0 \to 0$, observe that $N_i^F \to \infty$. Thus, the date at which second-mover adoption of a high-quality innovation converges in distribution to τ_G as $\alpha_0 \to 0$. Because $n_I^L(\alpha) > n_C^L(\alpha)$ for all $\alpha \in (0, \alpha_I^F)$, it follows that $\tau_I \mid_{\tilde{q}=q_H} >_{FOSD} \tau_C \mid_{\tilde{q}=q_H}$, as desired.

Now fix $\alpha_0 \in (0,1)$ and let $\lambda_B/\lambda_G \to \infty$. Observe that $\alpha_i^F \to 1$, so it is without loss to take $\alpha_0 < \alpha_I^F$. Conditional on no news, the incumbent adopts the innovation at a later date than does the challenger iff

$$\int_{\alpha_0}^{\alpha_I^F} \frac{n_I^L(\alpha) - n_C^L(\alpha)}{n_I^L(\alpha) n_C^L(\alpha) \alpha (1-\alpha)} d\alpha > \int_{\alpha_I^F}^{\alpha_C^F} \frac{1}{n_I^L(\alpha)\alpha(1-\alpha)} d\alpha.$$

As $\lambda_B/\lambda_G \to \infty$, the left-hand side term grows without bound, while the right-hand side term remains bounded above by

$$\frac{1}{n_I^L(0)}\left[\log\left(\Phi(\alpha_I^F)\right) - \log\left(\Phi(\alpha_C^F)\right)\right],$$

which is constant in λ_B/λ_G. Thus, $T_I > T_C$ holds whenever λ_B/λ_G is sufficiently large. When combined with our earlier comparison of τ_G, we conclude that $\tau_I \mid_{\tilde{q}=q_H} >_{FOSD} \tau_C \mid_{\tilde{q}=q_H}$, as desired. □

Proof of Theorem 2. Begin by noting that strict convexity of $\Pi(\cdot)$ guarantees that $L_i(\alpha) = 0 \implies \Pi(\Delta_i + \delta(\alpha;\tau_j)) > 0$. Thus, $\alpha_i^L < \hat{\alpha}_j$ where $\hat{\alpha}_j$ is defined by the equation $\delta(\hat{\alpha}_j;\tau_j) = 0$. From this, we immediately conclude that $\alpha_i^L < \alpha_j^F$. We know from previous results that $\alpha_I^F < \alpha_C^F$. Thus, both firms are unwilling to wait as second-movers if $\alpha_0 > \alpha_C^F$, but both are willing to adopt as first-movers. So the unique outcome in this case is simultaneous adoption. If $\alpha_0 \in (\alpha_I^F, \alpha_C^F)$, then both firms are willing to adopt as first-movers but only the challenger is willing to wait as a second-mover. So the unique outcome is initial adoption by the incumbent and waiting by the challenger.

To finish the characterization of initial adoption, we must compare α_I^L to α_C^L. To do this, it is helpful to first compute

$$\begin{aligned} L_i(\alpha) &= \theta_G(\alpha;\alpha_j^F)\left[\Pi(\Delta_i + x) - \Pi(\Delta_i)\right] + \theta_B(\alpha;\alpha_j^F)\left[\Pi(\Delta_i + q_L) - \Pi(\Delta_i)\right] \\ &= L_i(\alpha;1) - \frac{1}{\alpha}\left(\frac{\Phi(\alpha_j^F)}{\Phi(\alpha)}\right)^{\frac{\lambda_G}{\lambda_B - \lambda_G}}\left[1 + \Phi(\alpha_j^F)\right], \end{aligned}$$

where $L_i(\alpha;1) \equiv \alpha\left[\Pi(\Delta_i + x) - \Pi(\Delta_i)\right] + (1-\alpha)\left[\Pi(\Delta_i + q_L) - \Pi(\Delta_i)\right]$. As λ_B/λ_G becomes sufficiently large, we have $\alpha_i^L, \alpha_j^F \to 1$. So the second-mover waiting effect becomes small relative to the market structure effect in first-mover incentives. Thus, $L_I(\alpha) < L_C(\alpha)$ holds whenever $L_I(\alpha;1) < L_C(\alpha;1)$, which is equivalent to $0 > \alpha x + (1-\alpha)q_L \approx \delta(\alpha;\tau_i)$. It follows that $L_I(\alpha)$ intersects zero at a higher belief than does $L_C(\alpha)$ whenever λ_B/λ_G – or, equivalently, $\alpha_C^L < \alpha_I^L$. For all $\alpha_0 \in (\alpha_C^L, \alpha_I^L)$, both firms are willing to wait as second-movers but only the challenger is willing to adopt as a first-mover. So the unique outcome is that the challenger adopts while the incumbent waits. IF $\alpha_0 < \alpha_C^L$, then neither firm is willing to adopt as a first-mover and both are willing to wait as second-movers. So the unique outcome is that neither firm adopts the innovation. For $\alpha_0 \in (\alpha_I^L, \alpha_I^F)$, both firms are willing to adopt as a first-mover or wait as a second-mover. So multiple equilibria exist. □

Proof of Lemma 2. Let $n_{1,t}$ denote Firm 1's market share at time t. Then flow total surplus at time t is

$$\begin{aligned} TS_t &= \int_0^{n_{1,t}} \left[q_{1,t} + h(1-\theta)\right] d\theta + \int_{n_{1,t}}^1 \left[q_{2,t} + h\theta\right] d\theta \\ &= q_{1,t}n_{1,t} + q_{2,t}(1-n_{1,t}) + h\left[\frac{1}{2} + n_{1,t}(1-n_{1,t})\right]. \end{aligned}$$

The first-order condition for maximizing TS_t with respect to $n_{1,t}$ is $q_{1,t} + h(1-n_{1,t}) - q_{2,t} - hn_{1,t} = 0$, which solves to give the statically-optimal market share for Firm i, namely $n_{i,S}(\Delta_{i,t}) = (h + \Delta_{i,t})/(2h)$. To derive $TS(q_{i,t}, q_{j,t})$, insert the solution for $n_{i,S}(\cdot)$ into the above expression for TS_t and simplify. □

Proof of Proposition 4. The derivation of $\alpha_{i,S}^F$ is analogous to the proof of Proposition 2. All that we must verify is that second-mover adoption is (i) optimal after good news arrival, and (ii) suboptimal after bad news arrival. Second-mover adoption after good news arrival is strictly optimal iff

$$TS(q_i, q_j - x) + q_H > TS(q_i + q_H, q_j) \iff \frac{(q_H - x)(2h - 2\Delta_i - q_H - x)}{4h}.$$

This inequality holds under Assumption 3. Second-mover adoption after bad news arrival is strictly suboptimal iff

$$TS(q_i + q_L, q_j) > TS(q_i, q_j - x) + q_L \iff \frac{(x - q_L)(2h - 2\Delta_i - q_L - x)}{4h}.$$

This inequality holds under Assumption 3. ☐

Proof of Proposition 5. Using the strong long-run average criterion, the planner's dynamically-optimal choice of n_i given beliefs α is that which solves

$$\max_{n_i} \left\{ \frac{(q_i + q(\alpha))n_i + q_j(1-n_i) + h\left(\frac{1}{2} + n_i(1-n_i)\right) - v^L_{i,S}(\alpha)}{n_i} \right\} \\ \iff \max_{n_i} \left\{ \frac{h/2 - q_j - v^L_{i,S}(\alpha)}{n_i} - hn_i \right\}. \tag{32}$$

The first-order condition defining the optimal choice of n_i is $-(h/2 - q_j - v^L_{i,S}(\alpha))/(n_i)^2 = 0$, which solves to give $n_i = \sqrt{(v^L_{i,S}(\alpha) - q_j - h/2)/h}$, as desired.

To establish the comparison between $n^L_{i,S}(\alpha)$ and $n^L_i(\alpha)$, we first establish that $n^{(1,0)}_{i,S}(\alpha) - n^{(1,0)}_i(\alpha)$ is single-crossing. It suffices to prove that

$$\frac{d}{d\alpha}\left[n^{(1,0)}_{i,S}(\alpha) - n^{(1,0)}_i(\alpha)\right] = \frac{\left(v^L_{i,S}\right)'(\alpha)}{2hn^{(1,0)}_{i,S}(\alpha)} - \frac{\left(v^L_i\right)'(\alpha)}{4hn^{(1,0)}_i(\alpha)} > 0$$

holds whenever $n_{i,S}^{(1,0)}(\alpha) = n_i^{(1,0)}(\alpha)$. Clearly, this condition holds whenever $2(\nu_{i,S}^L)'(\alpha) - \left(\nu_i^L\right)'(\alpha) > 0$. For λ_B/λ_G sufficiently large, we have $(\nu_{i,S}^L)(\alpha) \approx \alpha\left(TS(q_i, q_j - x) + q_H\right) + (1-\alpha)TS(q_i + q_L, q_j)$ and $\nu_i^L(\alpha) \approx \alpha\Pi(\Delta_i + x) + (1-\alpha)\Pi(\Delta_i + q_L)$. Both $\nu_{i,S}^L(\alpha)$ and $\nu_i^L(\alpha)$ are continuously differentiable over $(0, \alpha_j^F)$. Therefore, $2(\nu_{i,S}^L)'(\alpha) - (\nu_i^L)'(\alpha)$ approximately equals

$$2\left(TS(q_i, q_j - x) + q_H - TS(q_i + q_L, q_j)\right) - (\Pi(\Delta_i + x) - \Pi(\Delta_i + q_L))$$

whenever λ_B/λ_G is sufficiently large. Under Assumption 3, the right-hand side of the above is positive. Therefore, $n_{i,S}^L(\alpha) - n_i^L(\alpha)$ is single-crossing over $(0, \alpha_j^F)$.

Direct verification shows that Assumption 1 guarantees $n_{i,S}^L(0) < n_i^L(0)$ and $n_{i,S}^L(\alpha_j^F) \approx n_{i,S}^L(\Delta_i + x) > \bar{n}(\Delta_i) = n_i^L(\alpha_j^F)$ whenever λ_B/λ_G is sufficiently large. Thus, Assumption 3 is sufficient to imply that $n_{i,S}^L(\alpha) - n_i^L(\alpha)$ crosses zero somewhere in $(0, \alpha_j^F)$.

Finally, to demonstrate that $n_{i,S}^L(\alpha) - n_i^L(\alpha)$ is increasing, it suffices to demonstrate that $n_{i,S}^{(1,0)}(0) - n_i^{(1,0)}(0) < 0$. This can be shown through direct simplification:

$$n_{i,S}^{(1,0)}(0) < n_i^{(1,0)}(0) \iff (q_i - q_j + q_L)[3h + 2(q_i - q_j + q_L)] < 0.$$

Thus, $n_{i,S}^{(1,0)}(\alpha) - n_i^{(1,0)}(\alpha)$ is strictly increasing over $(0, \alpha_j^F)$. □

Proof of Proposition 6. By previous results, we know that $\theta_G(\alpha; \alpha_{j,S}^F)$ and $\theta_B(\alpha; \alpha_{j,S}^F)$ are strictly concave and strictly convex, respectively. Direct calculations show that $TS(q_i, q_j - x) < TS(q_i, q_j) < TS(q_i + q_L, q_j) - q_L$ holds under Assumption 3. Thus, $L_{i,S}(\alpha)$ is strictly convex over $(0, \alpha^F_{j,S})$.

As $\alpha \to 0$, we have $L_{I,S}(\alpha) - L_{C,S}(\alpha) \to TS(q_i + q_L, q_j) - TS(q_i, q_j + q_L) = q_L \Delta_I / h < 0$. As $\alpha \to \alpha^F_{I,S}$, we have $L_{C,S}(\alpha) = 0$ because Firm I adopts the innovation as a second-mover almost instantaneously. However, we still have $L_{I,S}(\alpha) > 0$ because $\alpha^F_{C,S} > \alpha^F_{I,S}$. Therefore, $L_{I,S}(\alpha) > L_{C,S}(\alpha)$ holds as we let $\alpha \to \alpha^F_I$. It follows that $L_{i,S}(\alpha) - L_{j,S}(\alpha)$ crosses zero exactly once (by strict convexity) between $\alpha = 0$ and $\alpha = \alpha^F_{I,S}$. To guarantee that $L_{i,S}(\alpha) - L_{j,S}(\alpha)$ crosses zero when each curve is positive, we let λ_B / λ_G become sufficiently large so that $\alpha^L_{i,S}, \alpha^F_{j,S} \to 1$. Then the approximation $L_{i,S}(\alpha) - L_{j,S}(\alpha) \approx \Delta_i(\alpha x + (1-\alpha) q_L)/h$ holds. Setting this difference equal to zero and solving for α yields the following crossing belief: $\hat{\alpha}_S = (x - q_L)/(-q_L)$. By definition, $\hat{\alpha}_S$ is such that the expected long-run change in first-mover quality advantage is zero when the second-mover waits for news. Provided that λ_B / λ_G is sufficiently large, we are guaranteed that both $\alpha^L_{I,S}$ and $\alpha^L_{C,S}$ are less than $\hat{\alpha}$ because total surplus is strictly convex. Therefore, we conclude that $L_{I,S}(\alpha) - L_{C,S}(\alpha)$ crosses zero when both $L_{I,S}(\alpha)$ and $L_{C,S}(\alpha)$ are positive, which implies that $\alpha^L_{I,S} > \alpha^L_{C,S}$ holds. The characterization of socially efficient initial adoption follows immediately from this comparison. In particular, the planner chooses α^L_S to be the solution to the equation $L_C(\alpha) = 0$.

Finally, to compare α^L_S and α^L_C, we again let λ_B / λ_G become large. In the limit as $\lambda_B / \lambda_G \to \infty$, it can be verified that the following relationship holds: $2L_{i,S}(\alpha) = 9L_i(\alpha) - (\alpha x + (1-\alpha) q_L) + 2\alpha(q_H - x)$. Because $\alpha^L_i < \hat{\alpha}$, it follows that $L_{i,S}(\alpha) > 0$ at $\alpha = \alpha^L_i$. Thus, $\alpha^L_S < \alpha^L_C$. $\square$

Proof of Lemma 4. Suppose the two firms merge to become a joint-profit maximizing monopolist. Given qualities q_1, q_2 and prices p_1, p_2, a consumer purchases good 1 iff

$$q_1 + h(1-\theta) - p_1 \geq \max\{q_2 + h\theta - p_2, 0\}$$

$$\iff \theta \leq \min\left\{\frac{(q_1-p_1)-(q_2-p_2)+h}{2h}, \frac{q_1-p_1+h}{h}\right\}.$$

Let this threshold be denoted by θ_1. Likewise, a consumer purchases good 2 iff

$$q_2 + h\theta - p_2 \geq \max\{q_1 + h(1-\theta) - p_1, 0\}$$

$$\iff \theta \geq \max\left\{\frac{(q_1-p_1)-(q_2-p_2)+h}{2h}, \frac{p_2-q_2}{h}\right\}.$$

Let this threshold be denoted by θ_2. It is straightforward to verify that $\theta_1 = ((q_1-p_1)-(q_2-p_2)+h)/(2h)$ holds iff $\theta_2 = ((q_1-p_1)-(q_2-p_2)+h)/(2h)$. Thus, there are two cases to consider:

Case 1: (Serve the entire market)

$$\theta_1 = \frac{(q_1-p_1)+(q_2-p_2)+h}{2h} = \theta_2.$$

Case 2: (Exclude the middle)

$$\theta_1 = \frac{q_1-p_1+h}{h} < \frac{p_2-q_2}{h} = \theta_2.$$

To prove that Case 2: (Exclude the middle) never arises, assume by way of contradiction that $\theta_1^* < \theta_2^*$ at the monopoly solution. Then p_1^* and p_2^* solve the following maximization problem:

$$\max_{p_1,p_2} \; p_1 \cdot \underbrace{\left(\frac{q_1 - p_1 + h}{h}\right)}_{\theta_1} + p_2 \cdot \underbrace{\left(1 - \frac{p_2 - q_2}{h}\right)}_{1-\theta_2} \qquad \text{subject to} \qquad \theta_1 \le \theta_2.$$

Conjecture that the inequality constraint is slack at the optimum. Then we can disregard it in the first-order conditions, and the optimal prices are easily computed as $p_1^* = (q_1 + h)/2$ and $p_2^* = (q_2 + h)/2$.

We must now verify that the constraint $\theta_1 \le \theta_2$ is satisfied at these prices. Plug p_1^* and p_2^* into the formulas for θ_1 and θ_2 to obtain

$$\theta_1^* = \frac{q_1 - p_1^* + h}{h} = \frac{q_1 + h}{2h} \qquad \text{and} \qquad \theta_2^* = \frac{p_2^* - q_2}{h} = \frac{h - q_2}{2h}.$$

We see $\theta_1^* \le \theta_2^*$ holds only if $q_1 + q_2 < 0$ which is ruled out by Assumption 4. Thus, the monopolist never excludes the middle.

Next, we wish to establish the monopolist continues to supply both products. Note the monopolist serves the entire market only if the marginal type, denoted $\bar{\theta}$, receives zero utility; otherwise, the monopolist could raise both p_1 and p_2 by some small amount $\varepsilon > 0$ to increase profits. Therefore, the marginal type $\bar{\theta}$ satisfies $q_1 + h(1-\bar{\theta}) - p_1 = q_2 + h\bar{\theta} - p_2 = 0$. The prices p_1 and p_2 can then be written as functions of $\bar{\theta}$, namely $p_1 = q_1 + h(1 - \bar{\theta})$

and $p_2 = q_2 + h\bar{\theta}$. Thus, the monopolist's profit-maximization problem is

$$\max_{\bar{\theta}} \ \left(q_1 + h(1-\bar{\theta})\right)\cdot\bar{\theta} + \left(q_2 + h\bar{\theta}\right)\cdot(1-\bar{\theta}) \quad \text{subject to} \quad 0 \le \bar{\theta} \le 1.$$

The optimal choice of marginal type is $\bar{\theta}^* = (2h + q_1 - q_2)/(4h)$, which is interior under Assumption 1. Thus, the monopolist supplies both products in the static optimum. Upon simplification, the monopolist's flow profit given q_1 and q_2 is $\Pi^M(q_1, q_2) = (4h(q_1+q_2) + (q_1-q_2)^2 + 4h^2)/(8h)$. □

Proof of Proposition 7. The argument is analogous to the proof of Proposition 4. □

Proof of Proposition 8. Suppose the monopolist sets p_i and p_j so as to give Firm i market share n_i. Then the flow profit maximizing choice of p_i and p_j are $p_i = q_i + h(1-n_i)$ and $p_j = q_j + hn_i$. Thus, the monopolist's optimal choice of n_i in Stage 2, with Firm i a first-mover, is

$$\max_{n_i} \left\{ \frac{[q_i + q(\alpha) + h(1-n_i)]n_i + [q_j + hn_i](1-n_i) - v^L_{i,M}(\alpha)}{n_i} \right\}.$$

Solving the first-order condition for n_i gives the desired solution for $n^L_{i,M}(\alpha)$.

To compare $n^L_i(\alpha)$, $n^L_{i,S}(\alpha)$, and $n^L_{i,M}(\alpha)$, observe the following limits as λ_B/λ_G grows large:

1. $v^L_i(\alpha) \to \Pi(\Delta_i) + \alpha[\Pi(\Delta_i + x) - \Pi(\Delta_i)] + (1-\alpha)[\Pi(\Delta_i + q_L) - \Pi(\Delta_i)]$.
2. $v^L_{i,S}(\alpha) \to TS(q_i, q_j) + q(\alpha) + \alpha\left[TS(q_i, q_j - x) - TS(q_i, q_j)\right]$
$$+(1-\alpha)\left[TS(q_i + q_L, q_j) - TS(q_i, q_j) - q_L\right].$$

3. $v^L_{i,M}(\alpha) \to \Pi_M(q_i, q_j) + q(\alpha) + \alpha\big[\Pi_M(q_i, q_j - x) - \Pi_M(q_i, q_j)\big]$
$$+(1-\alpha)\big[\Pi_M(q_i + q_L, q_j) - \Pi_M(q_i, q_j)\big].$$

Proceeding in a similar manner to the proof of Proposition 5, it can be shown that $n^L_{i,S}(\alpha) - n^L_{i,M}(\alpha)$ and $n^L_{i,M}(\alpha) - n^L_i(\alpha)$ are both strictly increasing and cross zero once. Thus, $n^L_{i,M}(\alpha)$ falls between $n^L_i(\alpha)$ and $n^L_{i,S}(\alpha)$ at both low and high beliefs. By continuity, there must also exist a nonempty interval where $n^L_i(\alpha)$ falls between $n^L_{i,M}(\alpha)$ and $n^L_{i,S}(\alpha)$.

To formally demonstrate this latter point, note that

$$\big[n^{(1,0)}_{i,S}(\alpha) - n^{(1,0)}_i(\alpha)\big] - \big[n^{(1,0)}_{i,S}(\alpha) - n^{(1,0)}_{i,M}(\alpha)\big] = n^{(1,0)}_{i,M}(\alpha) - n^{(1,0)}_i(\alpha).$$

Let $\bar{\alpha}_S$ be the solution to $n^{(1,0)}_{i,S}(\alpha) - n^{(1,0)}_i(\alpha) = 0$; and let $\bar{\alpha}_M$ equal the solution to $n^{(1,0)}_{i,M}(\alpha) - n^{(1,0)}_i(\alpha) = 0$. Then examine two cases:

1. If $\bar{\alpha}_S < \bar{\alpha}_M$, then $n^{(1,0)}_i(\alpha) < n^{(1,0)}_{i,S}(\alpha) < n^{(1,0)}_{i,M}(\alpha)$ holds for all $\alpha \in (\bar{\alpha}_S, \bar{\alpha}_M)$. Because $n^{(1,0)}_{i,M}(\alpha) - n^{(1,0)}_{i,S}(\alpha)$ decreasing and $n^{(1,0)}_{i,S}(\alpha) - n^{(1,0)}_i(\alpha)$ is increasing, there is a range of beliefs in $(\bar{\alpha}_S, \bar{\alpha}_M)$ over which $0 < n^{(1,0)}_{i,S}(\alpha) - n^{(1,0)}_i(\alpha) < n^{(1,0)}_{i,M}(\alpha) - n^{(1,0)}_{i,S}(\alpha)$. Over this range, duopoly experimentation is more efficient.
2. If $\bar{\alpha}_M < \bar{\alpha}_S$, then $n^{(1,0)}_{i,M}(\alpha) < n^{(1,0)}_{i,S}(\alpha) < n^{(1,0)}_i(\alpha)$ holds for all $\alpha \in (\bar{\alpha}_M, \bar{\alpha}_S)$. Because $n^{(1,0)}_{i,S}(\alpha) - n^{(1,0)}_{i,M}(\alpha)$ is increasing and $n^{(1,0)}_i(\alpha) - n^{(1,0)}_{i,S}(\alpha)$ is decreasing, there is a range of beliefs in $(\bar{\alpha}_M, \bar{\alpha}_S)$ over which $n^{(1,0)}_{i,M}(\alpha) - n^{(1,0)}_{i,S}(\alpha) < n^{(1,0)}_i(\alpha) - n^{(1,0)}_{i,S}(\alpha) < 0$. Over this range, duopoly experimentation is more efficient.

In both cases, duopoly learning is most efficient for an interior range of α, as desired. □

Proof of Proposition 9. The argument is analogous to the proof of Proposition 6. □

Chapter 3

The Role of Market Structure in Competitive Experimentation

3.1 Introduction

In 1948, the Polaroid Corporation invented the first instant camera: a five-pound device with the ability to automatically develop low-quality sepia photographs. By the late 1970s, Polaroid cemented itself as a household name through unparalleled rates of innovation building upon this original breakthrough: monochrome film in 1950, color film in 1963, "no waste" instant film in 1972, autofocus lenses in 1978. Polaroid maintained this dominance in innovation through the 1980s, as it developed substantial capabilities in digital imaging and microelectronics through ambitious research programs that prioritized innovation at the expense of any notion of performance. "*Do not undertake the program unless the goal is manifestly important and its achievement nearly impossible*," wrote Polaroid founder, Edwin Land, in the 1980 Annual Report Letter to Shareholders.[1]

Despite enjoying a substantial technological lead in the "race to digital" compared to rivals (Kodak, Fuji, etc.), Polaroid began slashing its R&D investment into digital imaging and microelectronics in the 1990s. Exploratory investments in fiber optics, solar cells, and disk drives were abandoned; development of Polaroid's latest digital camera, the PDC-

---

[1] The information provided in this discussion can be found in Tripsas and Gavetti (2000).

300, was outsourced; and the Polaroid Microelectronics Lab, where most of Polaroid's basic research into microelectronics was conducted, was sold off. The purpose of this sudden divestiture: to place greater focus on existing products. "*[W]e have to focus on what value added we provide that's unique… substitute technologies such as inkjet or thermal technologies are interesting, but they're not here yet*," said then-CEO, Gary DiCamillo, to justify the company's decision (Rosenbloom and Pruyne, 1997).

Polaroid's ill-fated choice to abandon digital technology can be understood as the decision to prioritize exploitation of existing technology at the expense of exploration of new technology.[2] The idea behind the trade-off dates back to Schumpeter (1934) and is intuitive: "*Both exploration and exploitation are essential for organizations, but they compete for scarce resources. As a result, organizations make explicit and implicit choices between the two*" (March, 1991, pg. 71). Despite the ubiquity of this trade-off in a variety of industrial contexts, the fundamental question of how competitive forces shape decisionmaking between these two activities has received comparatively little attention. The reason is that innovation incentives are simultaneously determined by rivalry in R&D as well as in the existing market, which introduces an interesting but complex endogeneity between competition and innovation incentives.

In this dissertation chapter, I develop a tractable framework to analyze how dynamic competition unfolds in industries where firms simultaneously balance both innovation and

---

[2] Indeed, the trade-off between exploration and exploitation appeared to be top of mind for many of Polaroid's top leaders. "*Can we be a down and dirty manufacturer at the same time as we're an innovator over here? Can you have two different philosophies running simultaneously in the company?*" explained one senior manager in an interview about the company's strategic challenges (Tripsas and Gavetti, 2000, pg 1155).

product market incentives. The model considers a patent race between two firms to develop a new technology. This new technology is known to be superior to firms' existing technology, but its feasibility is uncertain. Experimentation is required to successfully develop the new technology, but this comes at the expense of investment into existing technology and, hence, current profitability. Importantly, firms are direct competitors in both R&D and the product market. Thus, firms must confront a strategic trade-off between exploration of the new technology (to ensure future competitiveness) and exploitation of the existing technology (to ensure current competitiveness).

In contrast to the existing literature, which typically assumes that firms are either potential entrants or exogenous competitors, my model allows firms to endogenously compete within a pre-existing market structure. This market structure is fully general, and my results are not tied to any particular specification. Instead, my analysis utilizies general properties of flow profits which are satisfied in most widely-used models of product market competition (e.g. Bertrand competition, Cournot competition, Hotelling competition). In particular, my model subsumes the classic patent race with hazard rate uncertainty (i.e. Choi, 1991) as a special case.

In the typical patent race, a firm strictly benefits from rival exit because fewer rivals means a lower chance of rival success (and, hence, a higher chance of own success). In my framework, however, a firm need not benefit from rival exit—in fact, rival exit can make a firm strictly worse off. The intuition behind this result, which I call *innovator disadvantage*, is that any firm that exits the patent race frees up productive resources to maximize competitiveness using the existing technology. As a result, firms that remains in the patent race are left at a competitive disadvantage in the market because their resources continue to be tied up in R&D.

Proposition 11 gives necessary and sufficient conditions for the patent race to *always* feature an innovator disadvantage over a non-empty interval of beliefs. These conditions do not require firms to be short-sighted or innovation to be slow-arriving. Instead, they depend solely on the extent to which investment in the existing technology enhances a firm's relative market power. In the special case where flow profits are exogenous, the patent race never features an innovator disadvantage. In more general market structures, however, innovator disadvantages are extremely common. Winner-take-all market structures, for example, always produce an innovator disadvantage for any possible parameter values.

Propositions 12 and 13 characterize the possible Markov-perfect equilibrium (MPE) outcomes of the patent race. They describe two basic classes of equilibria, the existence of which depends on the relative strength of two endogenous forces, which I call the *racing effect* and the *competition effect.* The racing effect measures the extent to whcih rivalry in the patent race stimulates own R&D incentives. As long as a one firm experiments with developing the new technology, the other firm feels pressure to experiment as well, because not doing so could result in loss of business if its rival succeeds. However, greater rival investment into the new technology implies greater gains to investment into the existing technology (because its rival is less competitive). Therefore, rivalry in the patent race may also generate a disincentive to invest in R&D. The strength of this disincentive is measured by the competition effect.

The dynamics of equilibrium investment in the patent race change considerably depending on whether the racing or competition effect dominates. If the racing effect dominates the competition effect, then MPE investment closely resemble that of the pure patent race. In particular, both firms *simultaneously* exit the patent race at a belief threshold that corresponds to an individually optimal, non-strategic exit threshold. But if the competition effect dominates the racing effect, then firms *sequentially* exit the patent race in every MPE because rivalry has the overall effect of discouraging investment. Furthermore, the timing of initial exit from the patent race is highly strategic. To decide the order of exit, firms decide the order and timing of sequential exit through either wars of attrition or preemption, depending on whether initial exit leads to an innovator advantage or to an innovator disadvantage. Thus, my model can explain in parsimonious fashion an important regularity (or, rather, irregularity) of the empirical work on R&D competition—namely, that the effect of rivalry on R&D incentives appears highly sensitive to industry specification (Gilbert, 2006b).

The analysis of MPE investment shows that strategic incentives in the patent race vary considerably across different market structures. To assess the welfare implications of this variation, I compute socially optimal investment within a classic linear-Cournot specification, which is flexible enough to induce every possible MPE outcome. I show that the total amount of R&D performed in every MPE is, on average, socially insufficient compared to the social optimum. However, the intensity of R&D (i.e. the length of time both firms stay in the patent race) may be socially excessive or socially insufficient. The intensity of R&D is socially excessive whenever the racing effect is sufficiently strong because firms are then motivated to perform R&D based on private considerations that are irrelevant for the social planner: namely, business-stealing. By comparison, the intensity of R&D is socially insufficient whenever the competition effect is sufficiently strong because firms are tempted to prematurely exit the patent race in order to maximize current competitiveness.

In the typical patent race, the competition effect is always zero; hence, symmetric firms exit the patent race simultaneously in every MPE. Thus, there is no distinction between the intensive and extensive margins of R&D, and a social planner can implement the social optimum with an appropriately-sized innovation prize. In general, however, prizes are insufficient to implement the social optimum due to the possibility—and perhaps social desirability—of sequential exit from the patent race. Therefore, my results support recent arguments that public policy should take a dynamic approach to subsidizing innovation (e.g. Langer and Lemoine, 2022). In particular, innovation subsidy programs must be carefully designed to not generate wasteful incentives for attrition or preemption.

To demonstrate the policy implications of my framework, I consider the effect of horizontal mergers on both innovation and consumer welfare. Specifically, I show that a merger's anticipated effect on innovation depends critically on the amount of consolidation that occurs (Proposition 18). On the one hand, if no consolidation occurs following the merger (i.e. there is common ownership), then firms' ability to coordinate production and investment substantially reduces the opportunity cost of R&D. In this case, post-merger innovation incentives exceed pre-merger incentives if and only if there is no Arrow replacement effect. On the other hand, if consolidation occurs (i.e. the merger is tantamount to the elimination of one firm), then post-merger innovation incentives are *always* lower than pre-merger incentives. The reason is that consolidation, along with enhanced market power, substantially raises the endogenous opportunity cost of performing R&D. Due to this "opportunity cost effect," the Arrow replacement effect is no longer necessary for competition authorities to presume that a merger will harm innovation.

The analysis of horizontal mergers highlight an important trade-off static and dynamic efficiency. In particular, a merger may generate static efficiency gains through the realization of synergies (e.g. economies of scale), but these exact synergies may reduce innovation incentives by exacerbating the replacement and opportunity cost effects. In Proposition 19, I show that static merger review is prone to overestimate the social benefits of merger-related synergies compared to a dynamic policy that considers both static and dynamic effects. This result stands in contrast to Nocke and Whinston's (2010) influential equivalence result and is precisely due to the endogeneity of both market structure and innovation incentives. Importantly, this result gives microfoundations to recent arguments made by US and European competition authorities based on an "innovation theory of harm," by which a statically pro-competitive merger can be deemed dynamically anti-competitive solely due to its negative effects on innovation.

The remainder of this dissertation chapter is organized as follows. In the next section, I discuss the contribution of my dissertation chapter in relation to the existing literature.

In Section 3.3, I describe my model and perform benchmark analysis. In Section 3.4, I provide a complete description of MPE and perform related comparative statics. I then consider welfare and antitrust implications in Sections 3.5 and 3.6, respectively. Finally, Section 3.7 concludes with a discussion of extensions and potential directions for future work.

3.2 Related Literature

This dissertation chapter contributes to the broad literature on R&D competition and to the growing literature on competitive experimentation.[3] In particular, this is the first dissertation chapter to analyze competitive experimentation within an endogenous market context.

The model I consider builds upon Choi (1991), who was the first to examine a patent race with hazard rate uncertainty. Choi's basic model has since been generalized to include, for example, variable effort choice (Malueg and Tsutsui, 1997), learning-by-doing (Doraszelski, 2003), multiple lines of R&D (Akcigit and Liu, 2016), and private information (Moscarini and Squintani, 2010). Although many of these models interpret firms as being direct competitors in an existing market structure, firms are assumed to have no tangible strategic objectives outside of the race itself.[4] In sharp contrast, my focus lies directly

---

[3] For an overview of the theoretical literature on R&D competition, see Reinganum (1989) and Gilbert (2006b). For an overview that focuses on antitrust considerations, see Federico et al. (2020) and Shapiro (2012).

[4] The pure patent race is well-suited to some settings, such as research tournaments (e.g. Taylor, 1995), where the context suggests no interaction outside of the race. In other situations, such as those where firms are interpreted as asymmetric competitors (e.g. Awaya and Krishna, 2021), the market structure requirements of the pure patent race are more difficult to interpret.

on the endogenous relationship between competition and R&D incentives. Therefore, my model can be used to analyze a richer set of comparative statics and counterfactuals related market structure, innovation, and welfare.

My analysis is thematically related to Awaya and Krishna (2021), who study a patent race between firms with asymmetric information processing abilities. In a main result, they provide conditions under which the most informed firm need not be the most profitable. This result is analogous to my innovator disadvantage result, which shows that the most innovative firm need not be the most profitable. The key difference between our two disadvantage results is that theirs originates from how asymmetric information affects observational learning, whereas mine originates from how asymmetric investment affects endogenous market competition.

The patent race with hazard rate uncertainty can be viewed as an early model of strategic experimentation with exponential bandits, first studied in Keller et al. (2005). As in Choi (1991), the exponential bandit framework has been extended in numerous directions; see Horner and Skrzypacz (2017) for a detailed survey. The strategic environment considered in much of this literature is one of pure information externalities: the actions of one player have no effect on the payoffs of another player except through the information communicated by the action. As a result, the primary focus of this literature is to characterize the extent to which free-riding incentives drive a wedge between privately-optimal and socially-optimal learning.

To study experimentation in more realistic contexts, such as R&D competition, Recent work has sought to introduce direct payoff externalities into strategic experimentation in the form of racing; examples include Akcigit and Liu (2016), Das and Klein (2020a,b), and Liu and Wong (2022). In these models, direct payoff externalities exist ex-post because the first player to achieve success using the risky action receives an additional prize, such as a patent. My model goes one step further by introducing direct payoff externalities in both ex-post payoffs (after success occurs) as well as interim payoffs (before success occurs).

To the best of my knowledge, the only other model of strategic experimentation to introduce direct payoff externalities through interim payoffs is Thomas (2021a). In Thomas's model, players have two actions, risky and safe, but the safe action is congestible in the sense that only one player can choose it at any given time. In equilibrium, the congestability of the safe arm creates a strategic incentive to preemptively take the safe action. This outcome is clearly analogous to the preemption equilibrium that obtains in my model whenever a sufficiently large innovator disadvantage exists. The key difference between our results is that preemption incentives arise in Thomas's model from exogenous congestion, whereas preemption incentives in my model arise from endogenous product market competition.

3.3 Preliminaries

3.3.1 Strategic Environment

The model considers a race between two symmetric firms, 1 and 2, to develop a new technology. Time is continuous, and the time horizon is infinite. At the start of the game, each firm competes in the market using an existing technology. Compared to this existing technology, the new technology is known to be superior. Therefore, the winner of the race receives a high continuation value V^+, while the loser of the race receives a lower continuation value V^-. For the sake of interpretation, it will be convenient to think of the values V^+ and V^- as being generated from patent protection. To facilitate this interpretation, let $\Pi^+ = \rho V^+$ ($\Pi^- = \rho V^-$) denote the flow value associated with winning (losing) the race, where $\rho > 0$ is the discount rate.

As long as a Firm i stays in the race, denoted $e_{i,t} = 1$, it must continue to allocate resources toward research and development (R&D). As a function of $(e_i, e_j) \in \{0, 1\}$, the flow payoff that Firm i receives is denoted $\Pi(e_i, e_j) - f e_i$, where $\Pi(e_i, e_j)$ is flow profit from the market, and $f > 0$ is a known constant. The idea behind why $\Pi(e_i, e_j)$ depends on (e_i, e_j) is that resources allocated toward R&D can also be used to improve a firm's competitiveness using the existing technology. Therefore, the true cost of R&D includes not just the exogenous flow cost, f, but also the endogenous opportunity cost, $\Pi(0, e_j) - \Pi(1, e_j)$. It is through this endogeneity that my model captures the general features of the trade-off between exploration and exploitation.[5]

For the sake of generality, I do not impose a particular market structure within the model. Rather, I maintain the following general assumptions:

Assumption 5. $\Pi(e_i, e_j)$ *is (i) non-increasing in* e_i*, and (ii) submodular in* (e_i, e_j)*.*

Observe that my model subsumes the pure patent race as a special case since $\Pi(e_i, e_j) = 0$ trivially satisfies Assumption 5. But more importantly, my model captures a wide variety of specifications in which $\Pi(e_i, e_j)$ is endogenously determined through strategic interaction. Below are just two classic examples:

Example 1 (Bertrand Competition)**.** Suppose that firms 1 and 2 are Bertrand competitors. Given prices $(p_{i,t}, p_{j,t})$, let the quantity demanded of Firm i's product at time t satisfy

$$D_i(p_{i,t}, p_{j,t}) = 1 \cdot \mathbb{1}[p_{i,t} < p_{j,t}] + \frac{1}{2} \cdot \mathbb{1}[p_{i,t} = p_{j,t}].$$

Suppose that Firm i can produce output at constant marginal cost c if all available resources are devoted to production (i.e. $e_i = 0$). But if a firm allocates some of its resources toward R&D (i.e. $e_i = 1$), then its lowest achievable marginal cost is $c + \phi$ for some $\phi > 0$. Then Firm i's (static) equilibrium profit given (e_i, e_j) is given by $\Pi(e_i, e_j) = \phi(1 - e_i)e_j$, which satisfies Assumption 5.

Example 2 (Cournot Competition)**.** Suppose that firms 1 and 2 are Cournot competitors in a market with inverse demand $P(Q_t) = a - bQ_t$, where $Q_t = q_{1,t} + q_{2,t}$ denotes industry

---

[5] The exact sources of the trade-off between exploration and exploitation are typically organizational frictions, such as financing constraints (Giebel and Kraft, 2019), limited managerial attention (e.g. Dosi, 1988), and time constraints (e.g. Radner and Rothschild, 1975).

output at time t. As in Example 1, suppose that Firm i's marginal cost takes the form $MC(e_i) = c + \phi e_i$ for some $\phi > 0$. For simplicity, let $2\phi < a + c$ so that marginal cost differences between the two firms are always non-drastic. Then Firm i's (static) equilibrium profit given (e_i, e_j) is

$$\Pi(e_i, e_j) = \frac{\big(a - c + \phi(e_j - 2e_i)\big)^2}{9b},$$

which satisfies Assumption 5.

As in Choi (1991), firms are initially uncertain about the ease of innovation. With probability $\beta_0 \in (0, 1)$, the new technology is feasible in the sense that it can be developed at a Poisson rate of $\lambda > 0$ as long as a firm performs R&D. With probability $1 - \beta_0$, however, the new technology is infeasible in the sense that development is impossible no matter how long a firm performs R&D. The decision to perform R&D amounts to strategic experimentation with an exponential bandit (c.f. Keller et al., 2005). For this reason, I use the terms R&D and experimentation interchangeably.

As the race unfolds, firms update beliefs about the new technology's feasibility in real time based on how much cumulative R&D has been performed. Conditional on no success, posterior beliefs at time t, denoted β_t, can be found using Bayes' rule:

$$\beta_t = \frac{\beta_0 e^{-\lambda \int_0^t (e_{1,s} + e_{2,s}) ds}}{\beta_0 e^{-\lambda \int_0^t (e_{1,s} + e_{2,s}) ds} + (1 - \beta_0)}. \tag{33}$$

Observe that (33) describes a decreasing function of t whenever $e_{1,t} + e_{2,t} > 0$. Thus, no success is bad news about the new technology's feasibility. Inverting the formula for β_t

yields the following expression for cumulative R&D as a function of initial and current beliefs:

$$N_t = \frac{1}{\lambda}\left[\log\left(\frac{1-\beta_t}{\beta_t}\right) - \log\left(\frac{1-\beta_0}{\beta_0}\right)\right]. \tag{34}$$

My solution concept is Markov-perfect equilibrium (MPE). There are well-known difficulties in the formal definition of continuous-time strategies whenever players take simultaneous actions (Fudenberg and Tirole, 1985; Simon and Stinchcombe, 1989). Therefore, I adopt the convention of defining continuous-time MPE in terms of limits of discrete-time MPE as the elapsed time between periods approaches zero. Consequently, we can imagine that each time t consists of two discrete subperiods, say t^- and t^+, in which exactly one firm moves. The probability that Firm 1 is the first mover in period t is described by a selection function $\xi : \mathbb{R}_+ \mapsto [0,1]$. Within this discretized game, an MPE exists for any period length $\Delta > 0$ (Maskin and Tirole, 2001). Furthermore, the limit of any sequence of discrete-time MPE outcomes as $\Delta \to 0$ is an MPE outcome of the continuous-time game (Simon and Stinchcombe, 1989).

Keeping in mind the above technicalities, we may nevertheless define MPE heuristically as follows. A Markovian strategy for Firm i is a pair of stopping times $\sigma_i = (\tau_i^1, \tau_i^2)$, where τ_i^k denotes the (possibly random) date at which Firm i exits the race given that beliefs are β and that $k \in \{1,2\}$ firms remain in the race. Given my formulation of continuous time, it is without loss of generality to assume that a tie-break occurs in the event that firms attempt to simultaneously exit the patent race. Indeed, Firm 1 is selected to exit before Firm 2 in the continuous-time game with probability $\alpha \equiv \lim_{\Delta \to 0} \xi(t)$. Likewise, Firm 2 is selected to exit before Firm 1 with probability $1-\alpha$. Whichever firm that is not selected to exit then has the opportunity to revise its exit decision. Finally, the Markovian strategy profile (σ_i, σ_j) constitutes an MPE if and only if Firm i's choice of σ_i maximizes its expected present-discounted payoff given (β, k, σ_j)), and vice versa for Firm j's choice of σ_j, given (β, k, σ_i).

3.3.2 Benchmark: The Pure Patent Race

Before analyzing the general model, it is helpful to consider a benchmark setting in which firms compete solely in terms of R&D. In particular, suppose that flow profit is exogenous and that business stealing effects of rival success are zero. This specification represents the standard model in the literature, and its assumptions guarantee that patent racing incentives are devoid of product market considerations. For this reason, I call this specification the pure patent race, and it is without loss to impose the normalization $\Pi(\cdot,\cdot) = \Pi^{-} = 0$.

Proposition 10. *There is a unique MPE of the pure patent race. In this equilibrium, both firms simultaneously exit the patent race at time* $\tau^* = \inf\{t \geq 0 : \beta_t \leq \beta^*\}$.

To prove Proposition 10, it suffices to compute Firm i's best response to the action $e_j \in \{0,1\}$ given beliefs $\beta \in (0,1)$. By the Principle of Optimality, the value Firm i receives from staying in the race (i.e. $e_i = 1$) satisfies the following equation:

$$\rho V_i(\beta) = -f + \beta\lambda[V^{+} - V_i(\beta)] - \beta\lambda e_j V_i(\beta) - (1+e_j)\lambda\beta(1-\beta)V_i'(\beta), \tag{35}$$

Hence, Firm i's (discounted) value in the patent race can be expressed as the sum of (i) the expected gains from own R&D, $-f + \beta\lambda[V^+ - V_i(\beta)]$, (ii) the expected losses from rival R&D, $-\beta\lambda e_j V_i(\beta)$, and (iii) the dynamic effect of learning, $-(1+e_j)\lambda\beta(1-\beta)V_i'(\beta)$. Learning enters negatively into (35) because posterior beliefs, conditional on no success, are decreasing over time.

As long as Firm i stays in the patent race, it pays the flow cost $f > 0$. Therefore, it cannot be optimal for Firm i to stay in the patent race for all $\beta \in (0,1)$. Instead, Firm i will exit once beliefs about innovation's feasibility are sufficiently pessimistic. Letting β^* denote this threshold belief, we have two boundary conditions: (i) $\rho V_i(\beta_1^*) = \Pi(0,0)$ (value matching) and (ii) $V_i'(\beta_1^*) = 0$ (smooth pasting). Inserting these boundary conditions into (35) yields the following solution:

$$0 = -f + \beta\lambda V^+ \iff \beta^* = \frac{f/\lambda}{V^+}. \tag{36}$$

Intuitively, the belief β^* is increasing in f/λ, the expected total cost of developing the new technology, and decreasing in V^+, the value of successful innovation. Observe also that β^* does not depend on $e_j \in \{0,1\}$ (this proves Proposition 10). Thus, Firm i's optimal R&D strategy in the pure patent race is effectively non-strategic: regardless of whether Firm i is alone or facing direct rivalry in the patent race, its incentive to perform R&D is the same. As I show in the next section, this is a non-generic feature of the pure patent race that critically depends on the assumption that investment is confined to the patent race.

3.4 Equilibrium Analysis

In this section, I analyze strategic R&D incentives when firms undertake investment in both the patent race *and* the product market. Proceeding via backward induction, I begin by considering a subgame in which one firm invests in the new technology as the *sole innovator*, while the other firm invests in the existing technology as the *sole non-innovator*.

3.4.1 The Sole Innovator's Problem

Suppose that only Firm i remains in the patent race. As the sole innovator, Firm i's problem is to decide when to stop investing in the new technology. To do this, Firm i compares the value of staying in the patent race, denoted $V_i(\beta)$, to the value of exit, denoted $V_{(0,0)} \equiv \Pi(0,0)/\rho$.

By the Principle of Optimality, the value of staying in the race at beliefs $\beta \in (0,1)$ satisfies the following equation:

$$\rho V_i(\beta) = \Pi(1,0) - f + \beta\lambda\left[V^+ - V_i(\beta)\right] - \lambda\beta(1-\beta)V_i'(\beta). \tag{37}$$

Similar to the pure patent race, $\rho V_i(\beta)$ can be expressed as the sum of (i) flow profit, $\Pi(1,0) - f$, (ii) the expected gains from R&D, $\beta\lambda[V^+ - V_i(\beta)]$, and (iii) the effect of learning, $-\lambda\beta(1-\beta)V_i'(\beta)$. The key difference is that the opportunity cost Firm i incurs to stay in the patent race now includes $\mathscr{C}(0) \equiv \Pi(0,0) - \Pi(1,0)$, the foregone value of alternative investment. Importantly, the magnitude of this additional cost is endogenously determined in the market and depends on, for example, the market power of each firm.

Example 3 (Hotelling Competition)**.** Suppose that Firms 1 and 2 are located at opposite ends of a Hotelling linear city, with Firm 1 located on the left ($\theta_1 = 0$) and Firm 2 located on the right ($\theta_2 = 1$). Consumers have unit demands and are uniformly distributed over [0,1]. Suppose that a consumer located at $\theta \in [0,1]$ receives utility $q_i - p_i - t(\theta - \theta_i)^2$ from purchasing Firm i's product, where q_i denotes quality, p_i denotes price, and $t > 0$ denotes the consumer's travel cost.

Each firm has zero marginal cost and can produce a high-value product ($v_i = v_H$) if it devotes all existing resources toward production (i.e., $e = 0$). But if a firm allocates some resources toward experimentation, then its highest achievable value is $v_H - \phi$ for some $\phi \in (0, v_H)$. Provided that $\phi < 3t$, Firm i's flow profit in any MPE is $\Pi(e_i, e_j) = (18t)^{-1}(3t + \phi(e_j - e_i))^2$, and the value of foregone investment into production is $\mathscr{C}(0) = (18t)^{-1}\phi(6t - \phi)$. Direct inspection reveals that $\mathscr{C}(0)$ is strictly increasing in t. Hence, Firm i perceives R&D as being more costly whenever it enjoys greater market power as a local monopolist.

In general, $\mathscr{C}(0)$ is positive whenever investment into the existing technology raises a firm's current profitability. This is true in the previous example because horizontal differentiation implies that firms enjoy local monopoly power. But if firms are undifferentiated Bertrand competitors, as in Example 1, then $\mathscr{C}(0)$ is zero because $e_j = 0$ implies that Firm i's market power always enjoys the same amount of market power—namely, none.

As long as Firm i stays in the patent race, it pays the flow cost $f > 0$ and receives flow profit $\Pi(1,0) < \Pi(0,0)$. Therefore, it cannot be optimal for Firm i to stay in the patent race for all $\beta \in (0,1)$. Instead, Firm i will exit the patent race whenever beliefs about innovation's feasibility become sufficiently pessimistic. Leting β_1^* denote this threshold belief, we have two boundary conditions: (i) $\rho V_i(\beta_1^*) = \Pi(0,0)$ (value matching) and (ii) $V_i'(\beta_1^*) = 0$ (smooth pasting). Inserting these boundary conditions into (37) and the general solution for $V_i(\beta)$ yields Lemma 7.

Lemma 7. *As the sole innovator, Firm i exits the race at time* $\tau_1^* = \inf\{t \geq 0 : \beta_t \leq \beta_1^*\}$, *where*

$$\beta_1^* = \frac{\rho}{\lambda}\left(\frac{f + \mathscr{C}(0)}{\Pi^+ - \Pi(0,0)}\right). \tag{38}$$

Intuitively, the belief β_1^* is increasing in $(f + \mathscr{C}(0))/\lambda$, the expected total cost of developing the new technology, and decreasing in $(\Pi^+ - \Pi(0,0))/\rho$, the value of successful innovation. Compared to a setting where $\Pi(e_i, e_j)$ is exogenous (i.e. constant), Firm i generally exits the patent race at higher beliefs it perceives R&D to be more costly. One key exception is the case of Bertrand competition, where $\Pi(0,0) = \Pi(1,0) = 0$. In this case, we have $\beta_1^* = (\rho/\lambda)(f/\Pi^+) = \beta^*$. One might be tempted to conjecture from this observation that strategic R&D incentives are the same in both settings. However, this conjecture turns is false because the strategic value of exit (to become the sole non-innovator) is substantially different.

3.4.2 Innovator Disadvantage

A key feature of my model is that investment affects not only the likelihood of success in the patent race, but also the intensity of competition in the product market. Thus, $(e_i, e_j) \in \{0,1\}^2$ affects each firm's expected value in MPE through direct and indirect payoff externalities. To describe how this affects strategic R&D incentives, we must compute the expected value each firm receives in MPE.

First, consider Firm i. Letting $\Phi(\beta) = (1-\beta)/\beta$ denote the odds-ratio of beliefs, the general solution to (37) takes the form $V_{(1,0)}(\beta) = \overline{V}_{(1,0)}(\beta) + \mathbb{C}_{(1,0)}(1-\beta)\Phi(\beta)^{\rho/\lambda}$, where

$$\overline{V}_{(1,0)}(\beta) = (1-\beta)\left(\frac{\Pi(1,0)-f}{\rho}\right) + \beta\left(\frac{\Pi(1,0)-f+\lambda V^{+}}{\rho+\lambda}\right),$$

and $\mathbb{C}_{(1,0)}$ is a constant of integration. To interpret this solution, note that $\overline{V}_{(1,0)}(\beta)$ is the expected value Firm i would receive if it chose to stay in the patent race indefinitely. This term clearly describes a lower bound on $V_{(1,0)}(\beta)$ since Firm i can exit the patent race at any time. Hence, the term in $\mathbb{C}_{(1,0)}(1-\beta)\Phi(\beta)^{\rho/\lambda}$ is positive and describes the option value Firm i receives from its ability to exit the patent race.

Next, consider Firm j. Given $\beta > \beta_1^*$, the expected value Firm j receives as the sole non-innovator satisfies the following equation:

$$\rho V_{(0,1)}(\beta) = \Pi(0,1) + \beta\lambda[V^{-} - V_{(0,1)}(\beta)] - \lambda\beta(1-\beta)V'_{(0,1)}(\beta). \tag{39}$$

The general solution to this equation takes the form $V_{(0,1)}(\beta) = \overline{V}_{(0,1)}(\beta) + \mathbb{C}_{(0,1)}(1 - \beta)\Phi(\beta)^{\rho/\lambda}$, where

$$\overline{V}_{(0,1)}(\beta) = (1-\beta)\left(\frac{\Pi(0,1)}{\rho}\right) + \beta\left(\frac{\Pi(0,1) + \lambda V^{-}}{\rho + \lambda}\right),$$

and $\mathbb{C}_{(0,1)}$ is a constant of integration. The interpretation of $V_{(0,1)}(\beta)$ is similar to that of $V_{(1,0)}(\beta)$ except for one key difference. Because Firm j does not choose $e_i \in \{0,1\}$, the option value it receives from Firm i's ability to exit the patent race may be positive or negative. Intuitively, as long as Firm i continues to (unsuccessfully) perform R&D, Firm j receives an indirect benefit from $e_i = 1$ in the form of enhanced market power, as measured by $\Pi(0,1) - \Pi(0,0)$. However, with R&D also comes the risk that Firm i successfully develops the innovation. This expected value of this risk is captured by the term $\beta\lambda[V_{(0,0)} - V^{-}]$.

In the pure patent race, both $\Pi(0,1) - \Pi(0,0)$ and $\beta\lambda[V_{(0,0)} - V^{-}]$ are zero. In this case, the exact solution for the sole non-innovator's value is $V_{(0,1)}(\beta) = 0$, which is always less than $V_{(1,0)}(\beta)$ over $(\beta_1^*, 1)$. In other words, the pure patent race always features a strict *innovator advantage* in MPE. As my next result demonstrates, this is a non-generic feature of the pure patent race that fails in a variety of market structures.

Proposition 11 (Innovator Disadvantage)**.** *There exists a non-empty interval of beliefs* $(\beta_1^*, \hat{\beta})$ *over which a strict innovator disadvantage exists (i.e.* $V_{(1,0)}(\beta) < V_{(0,1)}(\beta)$*) when-*

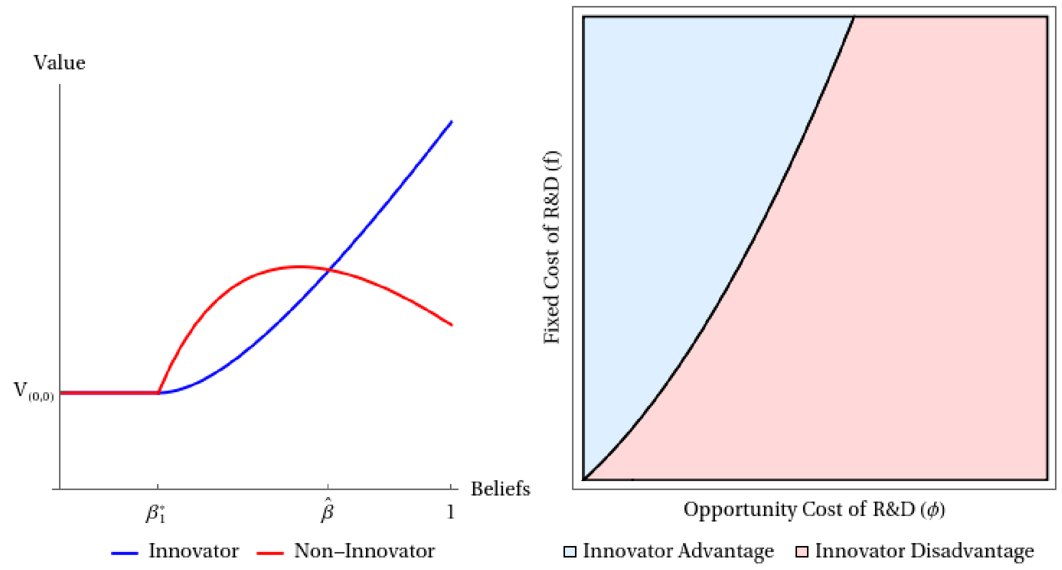

Figure 6: Innovator Disadvantage (Cournot Competition)

ever the following condition holds:

$$\Pi(0,1) - \Pi(0,0) \geq \kappa \equiv \left(\frac{\Pi(0,0) - \Pi^-}{\Pi^+ - \Pi(0,0)}\right)(f + \mathscr{C}(0)). \tag{40}$$

Proposition 11 demonstrates that the sole non-innovator need not just receive a positive benefit from rival investment in the patent race; it may very well receive a larger benefit than the sole innovator itself. The intuition for this result is similar to that of a joint profit effect. In particular, an innovator disadvantage will exist in MPE whenever asymmetric investment into the existing technology sufficiently relaxes product market competition, which disproportionately favors the sole non-innovator since $\Pi(0,1) - \Pi(0,0) > \Pi(0,0) - \Pi(1,0)$. Importantly, the existence of an innovator disadvantage does not require that firms be relatively impatient or that successful innovation be relatively slow. Instead, the existence of an innovator disadvantage depends exclusively on the underlying market structure in which firm's compete.

Example 4 (Bertrand Competition)**.** Suppose that $\Pi(e_i, e_j) = \phi(1-e_i)e_j$, as in Example 1. Then there always exists an innovator disadvantage over a nonempty interval of beliefs $(\beta_1^*, \hat{\beta})$ since $\Pi(0,1)-\Pi(0,0) = \phi > 0 = \kappa$. Intuitively, the reason that innovator disadvantages always arise in Bertrand competition is that the sole non-innovator is a de facto monopolist as long as the sole innovator continues to perform R&D and is unsuccessful.

Example 5 (Cournot Competition)**.** Consider the Cournot specification described in Example 2. Assuming that innovation is non-drastic (i.e. $2c < a + c^+$), the existence of an innovator disadvantage depends on the relative magnitudes of ϕ and f; in particular,

$$\Pi(0,1)-\Pi(0,0) > \kappa \iff \frac{4(a-c)\phi(c-c^+ + \phi)}{3b(2a-3c+c^+)} > f.$$

In words, an innovator disadvantage exists if and only if the non-innovator enjoys a sufficiently large marginal cost advantage over the sole innovator in the product market.

The possibility of an innovator disadvantage gives us a first glimpse at the novel strategic considerations that product market incentives generate. In the pure patent race, firms always would like to face fewer rivals in the patent race. However, this conclusion need not hold in markets with a significant innovator disadvantage. In such environments, Firm i can be made *worse-off* when its rival exits the patent race because this rival can then allocate its resources toward maximizing competitiveness in the product market. To avoid becoming the sole innovator, and suffering from an innovator disadvantage, firms may possess strategic incentives to underinvest in R&D in order to become the sole non-innovator and benefit from enhanced market power while their rival tries their luck at developing the possibly-infeasible new technology.

3.4.3 Strategic Exit: Attrition vs Preemption

Now suppose that both firms are still in the race. Our goal is to characterize the effect of rivalry in the patent race (i.e. $e_j = 1$) on Firm i's incentive to perform R&D. To do this, it is helpful to first consider a relaxed best response problem where Firm i may freely exit and re-enter the patent race. Letting $V_i(\beta)$ denote Firm i's expected value, we can use the Principle of Optimality to express this value recursively:

$$\rho V_i(\beta) = \max_{e_i \in \{0,1\}} \left\{ \begin{aligned} &\Pi(e_i, e_j) - f e_i + \beta\lambda e_i[V^+ - V_i(\beta)] \\ &\qquad + \beta\lambda e_j[V^- - V_i(\beta)] - (e_i + e_j)\lambda\beta(1-\beta)V_i'(\beta) \end{aligned} \right\}. \tag{41}$$

Inspect the terms on the right hand side of this equation, to see that $e_i = 1$ is a best response given (β, e_j) if and only if $\beta\lambda[V^+ - V_i(\beta) - (1-\beta)V_i'(\beta)] \geq f + \mathscr{C}(1)$. Inserting this condition back into the (41) yields an equivalent condition in terms of the value $V_i(\beta)$, namely that $e_i = 1$ is a best response give (β, e_j) if and only if $\rho V_i(\beta) \geq \mathscr{D}(\beta \mid e_j)$, where

$$\mathscr{D}(\beta \mid e_j) \equiv \Pi(0, e_j) + e_j\left[f + \mathscr{C}(e_j) - \beta\lambda(V^+ - V^-)\right]. \tag{42}$$

Holding all else equal, a higher value of $\mathscr{D}(\beta \mid e_j)$ means that Firm i chooses $e_i = 1$

over a smaller subset of beliefs. For this reason, we can think of $\mathscr{D}(\beta \mid e_j)$ as measuring that Firm i disincentive to perform R&D. Therefore, a natural question to ask is: how does $\mathscr{D}(\beta, e_j)$ depend on e_j? To answer this question, first note that $\mathscr{D}(\beta,1)-\mathscr{D}(\beta,0)$ is strictly decreasing in β. Thus, a necessary and sufficient condition for $\mathscr{D}(\beta,1)<\mathscr{D}(\beta,0)$ to hold for all $\beta \in (\beta_1^*,1)$ is

$$D(\beta_1^*,1) \leq D(\beta_1^*,0) \iff \underbrace{\beta_1^*\lambda(V^+-V^-)-f}_{\text{Racing Effect (RE)}} \geq \underbrace{2\Pi(0,1)-\Pi(1,1)-\Pi(0,0)}_{\text{Competition Effect (CE)}}. \tag{43}$$

From this, we see that rivalry in the patent race affects Firm i's incentive to perform R&D through two channels. First, there is a *racing effect* (RE) which measures the stimulative effects of $e_j = 1$ on Firm i's incentive to perform R&D. Intuitively, the racing effect measures the strength of business stealing incentives in the patent race; indeed, RE is positive if and only if $V^- < V_{(0,0)}$. Second, there is a *competition effect* (CE) that measures the strength of business stealing incentives in the product market. Indeed, as Firm j invests more resources into the patent race, it necessarily invests less into the current product market. This raises Firm i's temptation to invest in the existing technology as the sole non-innovator.

We can now precisely explain why simultaneous exit occurs at β_1^* in the pure patent race. In this setting, $V^- = V_{(0,0)} = 0$ implies that $RE = 0$. Thus, rivalry in the patent race does not make Firm i any more willing to perform R&D. Additionally, $\Pi(e_i,e_j)=0$ implies that $CE = 0$. In other words, rivalry in the patent race does not make Firm i any less willing to perform R&D. Combining these two observations, we see that Firm i's optimal

R&D strategy is the same regardless of the actions of Firm j. As my next result shows, this feature of the pure patent race carries over into any market specification in which RE dominates CE.

Proposition 12 (Racing MPE)**.** *Suppose that $RE \geq CE$. Then there is a unique MPE of the patent race, namely $(\tau_i^1, \tau_i^2) = (\tau_1^*, \tau_1^*)$ for $i \in \{1,2\}$. In other words, both firms simultaneously exit the patent race as soon as beliefs cross β_1^*.*

Whenever the RE dominates CE, the dynamics of MPE investment are identical to the pure patent race. Based on the intuition we just developed for RE and CE, it is not surprising that $RE = CE$ implies that MPE investment is the same as in the pure patent. Less obvious, however, is why the Proposition 12 extends to the case where $RE > CE$. To provide intuition for this result, suppose that racing incentives caused both firms to stay in the race until beliefs crossed $\tilde{\beta} < \beta_1^*$. By definition, the only reason that both firms perform R&D at $\beta \in (\tilde{\beta}, \beta_1^*)$ is to prevent their rival from being the first to develop the innovation. Over this interval of beliefs, Firm i strictly prefers that Firm j exit the patent race. But if this is true, then Firm i has an incentive to exit the race because then Firm j would immediately exit the race as the sole innovator.

In general, $RE \geq CE$ holds whenever the strategic benefits of investment into the existing technology are small (e.g. the pure patent race). In market structures where the strategic benefits of investment into the existing technology are large (e.g. winner-take-all markets), the dynamics of MPE investment change considerably. To describe these changes, it is helpful to define $\beta_2^* \equiv \inf\{\beta \in (\beta_1^*, 1) : \rho V_{(0,1)}(\beta) \geq \mathscr{D}(\beta, 1)\}$ as the lowest belief at which

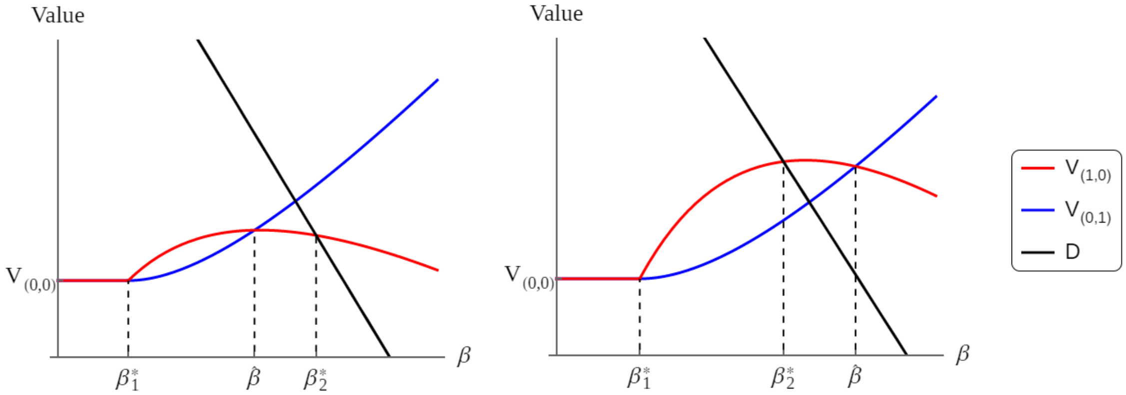

Attrition MPE

Preemption MPE

Figure 7: Exit Dynamics (Attrition vs Preemption)

Firm i would stay in the patent race if Firm j were expected to stay in the race.

Proposition 13 (Competition MPE)**.** *Suppose that $CE > RE$. Then there are two possibilities:*

1. *(*Attrition*) If there is an innovator advantage at β_2^*, then there are three possible MPE. In two of these equilibria, one firm exits the race as soon as beliefs cross β_2^*, while the other stays in the race until beliefs cross β_1^*. In the third equilibrium, both firms stay in the race with probability until beliefs cross β_2^*, then continuous randomization occurs. During this randomization, each firm exits with Poisson intensity*

$$\mu(\beta_t) = \frac{f + \Pi(0,0) + CE - \beta_t \lambda (V^+ - V^-) - \rho V_{(0,1)}(\beta_t)}{V_{(1,0)}(\beta_t) - V_{(0,1)}(\beta_t)}.$$

2. *(*Preemption*) If there is an innovator disadvantage at β_2^*, then there is an essentially unique MPE. In this equilibrium, one firm exits the race as soon as beliefs cross $\hat{\beta}$, while the other stays in until beliefs cross β_1^*.*

Whenever $CE > RE$, rivalry in the patent race decreases each firm's overall incentive to perform R&D. Firms sequentially exit the patent race in any MPE, and the timing of initial exit becomes highly strategic. To illustrate these strategic incentives, suppose that $CE > RE$ and suppose that both firms stay in until beliefs reach β_2^*. By definition, β_2^* is the belief at which neither firm is willing to continue in the race if its rival is expected to stay in. At the same time, $\beta_2^* > \beta_1^*$ (a consequence of $CE > RE$) implies that both firms are willing to stay in the race if their rival is expected to exit.

If there is an innovator advantage at β_2^*, then firms resolve this coordination problem by entering into a war of attrition. Noting the strategic equivalence of the patent race over $(\hat{\beta}, \beta_2^*)$ with the generalized war of attrition analyzed in Hendricks et al. (1988), we conclude that exactly three possible MPE exist. In two of these equilibria, one firm immediately exits the race. In the third, both firms randomize in hopes that their rival exits before they do. Staying in the patent race past β_2^* is costly to both firms; nevertheless, each is willing to continue as long as there is an innovator advantage. As time goes on, the value of winning the war of attrition, namely $V_{(1,0)}(\beta) - V_{(0,1)}(\beta)$, decreases. Thus, the rate at which firms exit the patent race is increasing over time; and as beliefs approach β_1^*, we have $\mu(\beta) \to \infty$. Therefore, the war of attrition almost-surely ends by time $\hat{\tau} = \inf\{t \geq 0 : \beta_t \leq \hat{\beta}\}$.

If there is an innovator disadvantage at β_2^*, then the war of attrition unravels. In this case, both firms strictly prefer to be the sole non-innovator, and one firm exits immediately at β_2^*. Anticipating this, the firm that expects to become the sole non-innovator at β_2^* has a strict incentive to preemptively exit the race slightly before time $\tau_2^* = \inf\{t \geq 0 : \beta_t \leq \beta_2^*\}$ in order to become the sole non-innovator. As in Fudenberg and Tirole (1985), this process continues for as long as there is an innovator disadvantage and ends at the belief $\hat{\beta}$, where rent equalization occurs. In other words, firms completely dissipate the strategic benefits of exit; see Figure 7.

3.4.4 Sufficient Conditions and Comparative Statics

As we showed in Proposition 11, the existence of an innovator disadvantage in MPE depends solely on market structure. The same is true about whether the timing of exit in MPE is sequential or simultaneous. Indeed, direct simplification shows that CE dominates RE if and only if

$$\Pi(0,1) - \Pi(0,0) > \left(\frac{\Pi(0,0) - \Pi^-}{\Pi^+ - \Pi(0,0)}\right)(f + \mathscr{C}(0)) - (\mathscr{C}(1) - \mathscr{C}(0)), \tag{44}$$

which is a weaker condition than (40) from Proposition 11 since $\mathscr{C}(1) \geq \mathscr{C}(0)$. Thus, an innovator disadvantage is not required for sequential exit to occur in MPE. This fact gives us an easy sufficient condition for attrition to occur in MPE—namely, that $\Pi(0,1) - \Pi(0,0)$ is large enough to imply that $CE > RE$ but not large enough to imply that $\hat{\beta} > \beta_1^*$.

Intuitively, we should expect that preemption is most likely to occur whenever the short-term benefits of investment into the existing technology are sufficiently large compared to the long-term benefits of successful innovation. My next result confirms this intuition by

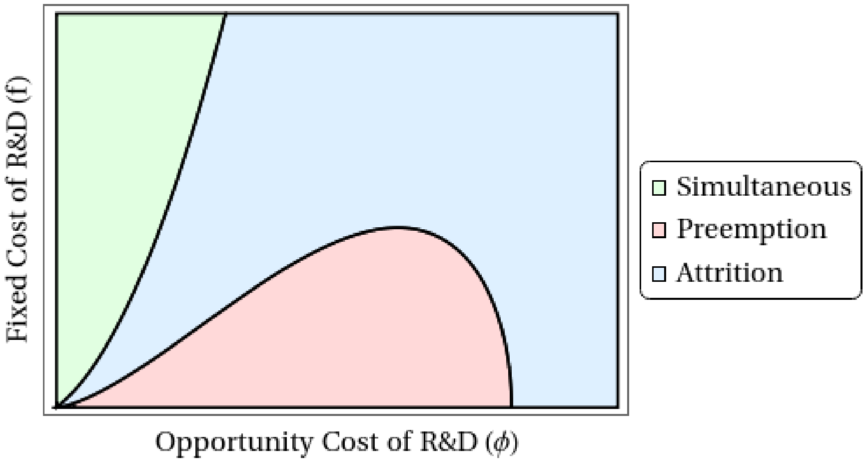

Figure 8: MPE Existence Regions (Cournot Competition)

examining the limit of MPE incentives as firms become relatively impatient.

Proposition 14 (Discounting and Preemption)**.** *Suppose that an innovator disadvantage exists and that* $\beta_0 > \beta_1^*$. *Then preemption is the unique MPE whenever firms are sufficiently impatient.*

Establishing sufficient conditions for preemption to occur in MPE in terms of parameters other than ρ is complex. To illustrate this complexity, consider the Cournot specification of Example 2. In this setting, the sole non-innovator enjoys a marginal cost advantage of size ϕ over the sole innovator, and an innovator disadvantage exists whenever ϕ is sufficiently large compared to f. Intuitively, the sole non-innovator's benefit from investing into the existing technology is strictly increasing in ϕ. From this observation, one might be tempted to think that preemption must occur in MPE whenever ϕ is sufficiently large. This logic turns out to be incorrect, however, because it does not account for effect of ϕ on the sole innovator's (endogenous) incentive to perform R&D.

Figure 8 depicts in (ϕ, f)-space where each type of MPE exists in the Cournot specification. This figure highlights the non-monotonic relationship between ϕ and MPE incentives for preemption. To develop intuition for this non-monotonicity, it is helpful to think about the direct and indirect effects through which an increase in ϕ affects the difference $V_{(0,1)}(\beta) - V_{(1,0)}(\beta)$. The direct effect of an increase in ϕ is to increase $\Pi(0,1) - \Pi(1,0)$. On the margin, this effect raises preemption incentives because the sole non-innovator is relatively more profitable. However, the indirect effect of an increase in ϕ is to increase β_1^* (i.e. decrease the sole innovator's R&D incentive). On the margin, this effect lowers preemption incentives because (i) the length of time over which the sole non-innovator receives $\Pi(0,1)$ is now shorter and (ii) the sole innovator less R&D at low beliefs (but still performs R&D at high beliefs).

3.5 Welfare Consequences

I now describe the welfare properties of MPE experimentation. To make precise statements about welfare, we must impose a particular market structure and compute total surplus. Because the Cournot industry specification is flexible enough to induce every possible MPE outcome at different parameter values, I maintain this specification throughout this analysis.

In the Cournot industry specification, total surplus may be computed as follows. Given a market equilibrium (P^*, Q^*), flow consumer surplus is

$$\int_0^{Q^*} (a - bQ - P^*)dQ = \left(a - \frac{b}{2}Q^*\right)Q^* - P^*Q^* = \frac{b}{2}(Q^*)^2.$$

Thus, pre-innovation flow consumer surplus given (e_i, e_j) is

$$CS(e_i, e_j) = \frac{b}{2}\left(q(e_i, e_j) + q(e_j, e_i)\right)^2 = \frac{\left(2(a-c) - \phi(e_i + e_j)\right)^2}{18b},$$

and post-innovation flow consumer surplus is

$$CS^+ = \frac{b}{2}\left(q^+ + q^-\right)^2 = \frac{(2a - c - c^+)^2}{18b}.$$

Likewise, pre-innovation flow total surplus given (e_i, e_j) is $TS(e_i, e_j) - f(e_i, e_j)$, where $TS(e_i, e_j) = CS(e_i, e_j) + \Pi(e_i, e_j) + \Pi(e_j, e_i)$, and post-innovation flow total surplus is $TS^+ = CS^+ + \Pi^+ - \Pi^-$. Note that I continue to maintain the assumption that marginal cost differences are non-drastic. This assumption is inconsequential for my analysis; its purpose is merely to simplify calculations.

3.5.1 Optimal Experimentation

To derive socially optimal R&D, consider the problem of a social planner who coordinates firms' investments to maximize present discounted total surplus. Given (β, e_1, e_2), the Bellman equation that defines social welfare, denoted $W(\beta)$, is

$$\rho W(\beta) = TS(e_i, e_j) - f(e_i + e_j) + \beta\lambda(e_i + e_j)[W^+ - W(\beta) - (1-\beta)W'(\beta)], \quad (45)$$

where $W^+ = TS^+/\rho$ denotes post-innovation welfare. Note that the planner does not distinguish between innovation by Firm i and innovation by Firm j because the two firms are symmetric; hence, the social value of innovation by each firm is the same.

As in the strategic problem, Firm i's social incentive to perform R&D depends on $e_j \in \{0,1\}$. To determine whether it is increasing or decreasing in e_j, let $W_k(\beta)$ denote social welfare at β given $k \in \{0,1,2\}$ firms in the race. By definition, $W_0(\beta) = TS(0,0)/\rho$ for all β. In this case, the indifference condition for optimal exit by Firm i given $e_j = 0$ is found by inserting the value matching condition $W_1(\beta) = TS(0,0)/\rho$ and smooth pasting condition $W_1'(\beta) = 0$ into the (45):

$$TS(0,0) = TS(1,0) - f + \beta\lambda\left[W^+ - \frac{TS(0,0)}{\rho}\right] \iff \beta = \beta_1^S \equiv \frac{\rho}{\lambda}\left(\frac{f + \mathscr{C}_S(0)}{TS^+ - TS(0,0)}\right),$$

where $\mathscr{C}_S(e_j) \equiv TS(0,e_j) - TS(1,e_j)$ denotes the endogenous social opportunity cost of R&D. The indifference condition for optimal exit by Firm i given $e_j = 1$ is similarly found by inserting $W_2(\beta) = W_1(\beta)$ and $W_2'(\beta) = W_1'(\beta)$ into (45), which yields $\rho W_1(\beta) = 2TS(0,1) - TS(1,1)$ upon simplification. This indifference condition is analogous to that which defines β_2^* in the strategic problem, except the social planner does not experience a racing effect.

In contrast to the strategic problem, where e_i and e_j could be complements or substitutes, the social planner always views e_i and e_j as substitutes. To see this, note that $\rho W_1(\beta) = TS(0,0)$ for all $\beta < \beta_1^S$ since the planner has Firm i exit immediately as the sole innovator. From this, we conclude that the equation $\rho W_1(\beta) = 2TS(0,1) - TS(1,1)$

admits a solution $\beta_2^S \geq \beta_1^S$ if and only if $TS(0,0) \leq 2TS(0,1) - TS(1,1)$. Since $TS(0,1) = TS(1,0)$ by symmetry, we see that $\beta_2^S \geq \beta_1^S$ whenever $TS(e_i, e_j)$ is submodular. Direct calculation shows that $TS(e_i, e_j)$ is always submodular since $TS(0,1) + TS(1,0) - TS(0,0) - TS(1,1) = \frac{7\phi^2}{9b}$ is non-negative (positive) for all $\phi \geq 0$ ($\phi > 0$). We conclude from this observation that the timing of exit in the optimal experimentation policy is sequential (simultaneous) whenever $\phi > 0$ ($\phi = 0$).

Proposition 15 (Optimal Experimentation)**.** *In the optimal experimentation policy, the last firm to exit the patent race does so when beliefs cross*

$$\beta_1^S = \frac{\rho}{\lambda}\left(\frac{f + \mathscr{C}_S(0)}{TS^+ - TS(0,0)}\right). \tag{46}$$

The first firm to exit the patent race does so when beliefs cross β_2^S, defined as the (unique) solution to the equation $W_1(\beta) = 2TS(0,1) - TS(0,0)$, where

$$W_1(\beta) = \overline{W}_1(\beta) + \frac{(1-\beta)\Phi(\beta)^{\frac{\rho}{\lambda}}}{(1-\beta_1^S)\Phi(\beta_1^S)^{\frac{\rho}{\lambda}}}\left(\frac{TS(0,0)}{\rho} - \overline{W}_1(\beta_1^S)\right) \tag{47}$$

and

$$\overline{W}_1(\beta) = (1-\beta)\left(\frac{TS(1,0) - f}{\rho}\right) + \beta\left(\frac{TS(1,0) - f + \lambda W^+}{\rho + \lambda}\right).$$

3.5.2 Equilibrium Inefficiency

In general, there are two possible types of equilibrium inefficiency. First, there can be inefficiency in the amount of R&D performed before all firms exit the patent race. This type of inefficiency is common across all MPE because the last firm to exit the patent race always does so at time $\tau_1^* = \inf\{t \geq 0 : \beta_t \leq \beta_1^*\}$. Therefore, total amount of R&D performed before all firms exit the race, in any MPE, is equal to

$$N_1^* = \int_0^{\tau_1^*} = (e_{1,t}^* + e_{2,t}^*)dt = \int_{\beta_1^*}^{\beta_0} \frac{d\beta}{\lambda\beta(1-\beta)} = \frac{1}{\lambda}\left[\log(\Phi(\beta_1^*)) - \log(\Phi(\beta_0))\right].$$

Second, there can be inefficiency in the amount of R&D performed before one firm exits the patent race. This type of inefficiency varies across MPE since exit may be simultaneous or sequential (attrition or preemption) and is thus indeterminate.

Despite this indeterminacy, we can obtain upper and lower bounds on the the intensity of MPE experimentation as follows. In every MPE, we know that initial exit occurs (w.p.1) by time $\hat{t} = \inf\{t \geq 0 : \beta_t \leq \hat{\beta}\}$. Thus, an upper bound on the intensity of R&D in any MPE is

$$\hat{N} = \int_0^{\hat{\tau}} (e_{1,t}^* + e_{2,t}^*)dt = \int_{\hat{\beta}}^{\beta_0} \frac{d\beta}{\lambda\beta(1-\beta)} = \frac{1}{\lambda}\left[\log(\Phi(\hat{\beta})) - \log(\Phi(\beta_0))\right].$$

If $\hat{\beta} > \beta_2^*$, then $\hat{N}$ also describes a lower bound on the intensity of MPE experimentation; otherwise, a lower bound on the intensity of experimentation is

$$N_2^* = \int_0^{\tau_2^*} (e_{1,t}^* + e_{2,t}^*)dt = \int_{\beta_2^*}^{\beta_0} \frac{d\beta}{\lambda\beta(1-\beta)} = \frac{1}{\lambda}\left[\log(\Phi(\beta_2^*)) - \log(\Phi(\beta_0))\right].$$

My next result characterizes the (in)efficiency of total R&D incentives.

Proposition 16. *For all $\phi \geq 0$, we have $\beta_1^* > \beta_1^S$. Thus, in any MPE, the total amount of R&D performed before all firms exit the patent race is less than the social optimum.*

Total R&D incentives are inefficient for two reasons: (i) failure to internalize the gains to successful innovation, and (ii) failure to internalize the benefits of asymmetric firm investment. Indeed, firms does not capture full social surplus generated from success since $\Pi^+ - \Pi(0,0) < TS^+ - TS(0,0)$, and the sole innovator perceives the opportunity cost of R&D as being higher than the social cost since $\mathscr{C}(0) > \mathscr{C}_S(0)$.

In terms of the intensity of R&D, the main source of inefficiency in MPE comes from either preemption or attrition incentives. As we saw in the previous subsection, the social planner does not perceive any difference between successful innovation by Firm i and successful innovation by Firm j. Hence, sequential exit always occurs at β_2^S. Sequential exit in MPE, by comparison, may occur at $\hat{\beta}$ in the case of preemption or at $\tilde{\beta} \in (\hat{\beta}, \beta_2^*]$ in the case of attrition. As my next result demonstrates, the (in)efficiency of MPE experimentation in terms of its intensity critically depends on the relative strength of business stealing incentives in the patent race which, in turn, depends on the level of patience of each firm.

Proposition 17. *Suppose that $CE > RE$ (so that exit is sequential). Then the intensity of MPE experimentation is (i) socially excessive (i.e. $\beta_2^* < \beta_2^S$) whenever firms are sufficiently patient and (ii) socially insufficient (i.e. $\hat{\beta} > \beta_2^S$) whenever firms are sufficiently impatient.*

The proof of Proposition 17 compares CE and $RE(\beta)$ at high and low discount rates. Because the social planner does not experience a racing effect, the threshold β_2^S is bounded away from zero for all $\rho > 0$. In contrast, each firm's racing incentive in MPE, namely

$RE(\beta) = \beta\lambda(V^+ - V^-) - f$ grows arbitrarily large as $\rho \to 0$. Consequently, β_2^* is arbitrarily small when firms are sufficiently patient. Conversely, if firms are relatively impatient, then the timing of sequential exit in MPE is primarily determined by the competition effect. In this case, firms are tempted to exit the patent race earlier than is socially optimal because they firms overvalue the benefits of enhanced market power compared to the social planner.

In the pure patent race, simultaneous exit occurs in MPE as well as in the social optimum since both $\Pi(e_i, e_j) = \Pi$ and $TS(e_i, e_j) = TS$ are constant. As a consequence, there is no distinction between the amount of experimentation that is performed in MPE and the intensity of experimentation performed in MPE. Moreover, it is possible to implement the social optimum with an innovation prize equal to $S = (TS^+ - TS) - (\Pi^+ - \Pi)$ whenever private incentives do not align with the social optimum. However, an innovation prize is not sufficient to achieve social efficiency when firms can strategically invest in the product market. The reason is that such a prize will correct total R&D incentives but may further distort the intensity of R&D by generating additional incentives for attrition or preemption. Instead, dynamic subsidies or taxes are required to achieve complete social efficiency.

To properly design the optimal dynamic subsidy or tax, the planner must consider whether the intensity of R&D is socially excessive or socially insufficient. If the intensity of MPE experimentation is socially insufficient, then the social planner can discourage exit via the imposition of a tax on non-experimenting firms. Such a tax is optimal so long as beliefs remain above β_2^S. After beliefs cross this threshold, the planner can remove the tax to induce sequential exit. If the intensity of MPE experimentation is socially excessive, then social efficiency can be implemented via subsidies on the existing technology. In particular, the social planner can offer a subsidy (or simply a lump sum payment) to the first non-experimenting firm that applies for it. By construction, this subsidy reduces the inefficiently high intensity of R&D with no effect on total R&D incentives since the subsidy applies only to the first firm to exit the race.

3.6 Antitrust Implications

A growing number of high-profile merger cases (e.g. Bayer/Monsanto, Dow/DuPont, Ilumina/PacBio, Sabre/Farelogix) have focused on concerns that merging parties may decrease innovation and, hence, long-term welfare.This concern appeared prominently in the Dow/DuPont merger case, where the European Commission articulated a view that mergers generally stifle innovation (Denicolo and Polo, 2018). In particular, the European Commission articulated an innovation theory of harm by which a merger with benign (or even positive) static effects can be regarded as anticompetitive from a dynamic perspective solely based on its anticipated effects on innovation. In this section, I demonstrate the policy implications of my framework with an application to merger review and a formal analysis of the innovation theory of harm.

To model the effects of mergers on innovation, the literature has traditionally focused on “product-to-pipeline” competition, where one firm has an existing product, and the other has a competing product in development, or “pipeline-to-pipeline” competition, where both firms have competing products in development (see Shapiro, 2011). Models to analyze mergers of each type have been applied to the pharmaceutical and health science industries (e.g. Celgene/Bristol-Myers Squibb, Ilumina/PacBio). However, these models have limited applicability to mergers in which firms simultaneously compete in both the product market and innovation (e.g. Dow/DuPont and Sabre/Farelogix). As a result, competition authorities have been forced to take more conservative paths to enforcement because "identifying acceptable divestiture remedies involves predicting the innovation consequences of altering the market structure" (Majure, 2022).

3.6.1 Mergers and Innovation

In traditional merger analysis, there are two main approaches to modeling a merger: coordination and consolidation (see Farrell and Shapiro, 1990; Whinston, 2007). In a merger with coordination, the merged entity—call it Firm M—retains the productive capabilities of both firms and coordinates production to maximize joint profit. In a merger with consolidation, Firm M scraps all overlapping productive capabilities and produces output as a single firm; in other words, consolidation is tantamount to the elimination of one firm. Absent synergies, the two approaches are considered to be equivalent in the analysis of mergers between Cournot competitors with constant marginal costs (McAfee and Williams, 1992).

This equivalence does not hold in my model because coordination and consolidation imply different innovation capabilities. In a merger with coordination, Firm M retains the innovation capabilities of both firms; thus, it chooses $e_M \in \{0, 1, 2\}$ in the patent race. In a merger with consolidation, Firm M retains the innovative capabilities of only one firm;

thus, it chooses $e_M \in \{0,1\}$. As my next result shows, this distinction implies substantially different post-merger innovation incentives.

Proposition 18 (Mergers and Innovation)**.** *Maintaining the Cournot specification from Example 2, suppose that $\phi > 0$. Then:*

1. *If the merger involves coordination, Firm M will perform R&D until beliefs cross*

$$\beta^M_{Coord} = \frac{\rho}{\lambda}\left(\frac{f}{\Pi^+_M - \Pi_M(0)}\right) = \frac{\rho}{\lambda}\left(\frac{4bf}{(c-c^+)(2a-c-c^+)}\right).$$

   In this case, Firm M always performs ***more*** *R&D than occurs in duopoly.*

2. *If the merger involves consolidation, Firm M will perform R&D until beliefs cross*

$$\beta^M_{Cons} = \frac{\rho}{\lambda}\left(\frac{f}{\Pi^+_M - \Pi_M(0)}\right) + \frac{\rho}{\lambda}\left(\frac{\phi}{c-c^+}\right)\left(\frac{2(a-c)-\phi}{2(a-c)+(c-c^+)}\right).$$

   In this case, Firm M always performs ***less*** *R&D than occurs in duopoly.*

In one of his seminal works, Arrow (1962) provided conditions under which the incentives to innovate are lower in a monopolistic market than in a competitive market. His analysis revolves around showing that, despite enjoying greater appropriability, a monopolist gains less from successful innovation than a competitive firm due to the cannibalization of existing profits. We can see this cannibalization effect at work in my model by setting $\phi = 0$. In this case, coordination and consolidation are equivalent from the standpoint of total R&D incentives (the intensity of R&D remains higher under coordination), and Firm M has stronger incentives to innovate than either of the pre-merger duopolists if and only if $\Pi_M^+ - \Pi_M(0) > \Pi^+ - \Pi(0,0)$, which simplifies to become $c - c+ < 2(a-c)/7$ in our Cournot specification. In other words, a merger leads to more R&D investment if and only if innovation is sufficiently incremental in nature.[6].

If firms can invest into maximizing current profitability in addition to R&D investment (i.e. $\phi > 0$), then monopoly innovation incentives change considerably. The first part of Proposition 18 shows that a merger with coordination always leads to more innovation. The intuition behind this result is the following. Because Firm M maintains the joint research and production capabilities of both firms, it is able to perform R&D and maximize current profitability. As a result, Firm M always performs partial experimentation (i.e. $e_M = 1$) before exiting the patent race entirely. During this partial experimentation regime, Firm M pays no cost to perform R&D besides f, which dramatically increases Firm M's willingness to stay in the patent race compared to the sole innovator in duopoly. Note that this result holds regardless of whether the $\Pi_M^+ - \Pi_M(0)$ is greater than or smaller than $\Pi^+ - \Pi(0,0)$. In other words, the Arrow replacement effect is no longer a sufficient to conclude that a merger harms innovation.

If consolidation occurs, then Firm M's innovation incentives are sharply reduced. The reason is that Firm M, like the sole innovator, must sacrifice short-term profitability in order to stay in the race. Compared to the sole innovator, however, it is much more costly for Firm M to perform R&D —that is, $\mathscr{C}_M > \mathscr{C}(0)$—because it has greater market power and, thus, can appropriate greater returns from investment into the existing technology. Due to

---

[6] Reinganum (1982) obtains a similar result in the context of an incumbent monopolist versus a potential entrant.

this *opportunity cost effect*, Firm M's incentives to innovate (as measured by total R&D performed before exit) are always lower than duopoly incentives. Importantly, this result holds regardless of $\Pi_M^+ - \Pi_M(0)$ is greater than or less than $\Pi^+ - \Pi(0,0)$. Combining this with the previous observation, we conclude that the Arrow replacement effect is neither sufficient nor necessary to conclude that a merger harms innovation. This result stands in contrast to Bourreau and Jullien (2018) who analyze a framework where mergers spur innovation precisely because they increase appropriability. A critical feature of their model is its lack of dynamics: innovation does not require that firms allocate resources toward R&D (instead of the product market) over an extended period of time.

As Proposition 18 shows, a distinctive advantage of my model is its ability to simultaneously consider both unilateral price effects and unilateral innovation effects. From an application standpoint, my results provide theoretical underpinnings to several recent merger cases. Consider, for example, the 2019 DOJ complaint challenging Sabre's acquisition of Farelogix. In this complaint, the DOJ alleged two separate theories of harm: (i) higher prices due to the elimination of head-to-head competition, and (ii) reduced incentives to invest and innovate next-generation technology. Furthermore, my model indicates that these two "separate" theories of harm are hardly separate. Instead, they are inextricably linked.

3.6.2 Mergers and Welfare

Horizontal merger review typically stresses the trade-off between market power and efficiencies (Farrell and Shapiro, 1990). To capture such efficiencies, suppose that a merger between Firms 1 and 2 realizes merger-related synergies that allow Firm M to produce at lower cost than either firm could achieve individually. In particular, let the marginal cost of Firm M be $(1-s)\times 100$ percent lower than the minimum achievable cost in duopoly, where $s \in [0,1]$ denotes the size of merger-related synergies. If Firm M is created via consolidation, then Firm M's pre-innovation marginal cost given $e \in \{0,1\}$ is $(1-s)(c+\phi e)$. If Firm M is created via coordination, then Firm M's pre-innovation marginal cost given $e \in \{0,1,2\}$ is $(1-s)c$ if $e \leq 2$ and $(1-s)(c+\phi)$ if $e=2$. Regardless of the level of consolidation, Firm M's post-innovation marginal cost is $(1-s)c^{+}$.

Suppose the two firms propose to merge before the start of the game. When this proposed merger is evaluated from a static perspective, the comparison between market power and synergies is relatively simple. Using a consumer welfare standard, the merger will be approved whenever the resulting monopoly consumer surplus, denoted $CS_M(0)$, exceeds duopoly total surplus, $CS(0,0)$. Direct calculation shows that ex-ante static consumer welfare improves as a result of the merger (with or without consolidation) if and only if $a < 4c$ and $s > \hat{s} \equiv (a-c)/3c$.

However, the dynamic relationship between synergies and welfare is more complex. The reason is that synergies affect reduce Firm M's endogenous innovation incentives, as the following corollary to Proposition 18 demonstrates.

Corollary 2. *Regardless of the level of consolidation, synergies always reduce the amount of R&D that Firm M performs before exiting the race.*

The intuition behind the above comparative static result is that synergies make Firm M more efficient at any level of investment, which reduces Firm M's gains to successful innovation. In other words, the difference $\Pi_M^+ - \Pi_M(0)$ is increasing in s. If the merger involves coordination, it follows immediately that Firm M's incentive to perform R&D, as measured by β_{Coord}^M, is increasing in s. If the merger involves coordination, however, the comparative static is less obvious because the endogenous opportunity cost of R&D, namely $\mathscr{C}_M = \Pi_M(0) - \Pi_M(1)$ is decreasing in s. Intuitively, this is because Firm M is more efficient regardless of its choice of investment, so it sacrifices less while performing R&D. Despite this beneficial effect of synergies on the opportunity cost of R&D, the proof of Corollary 2 shows that this effect is always small compared to the disincentive to invest caused by lower gains to successful innovation.

In light of this result, we can see the subtleties of assessing a merger's dynamic welfare effects. On the one hand, greater synergies imply higher consumer surplus in both the pre- and post-innovation monopoly market structures. From this perspective, greater synergies would suggest the merger is more likely to be pro-competitive. However, it is exactly these synergies that reduce also innovation incentives by reducing Firm M's gains from successful innovation. From this perspective, greater synergies would suggest the merger is more likely to be anti-competitive. By definition, a static merger review policy designed to maximize static consumer surplus does not account for this latter effect of merger-related synergies. As a result, there will be general disagreement between statically optimal and dynamically optimal merger review.

To illustrate the disagreement between static and dynamic merger enforcement, consider two contrasting merger review policies. The first is a static policy that approves a merger if and only if ex-ante consumer surplus increases–that is, $CS_M(0) > CS(0,0)$. As shown earlier, this static policy approves a merger, given $a < 4c$, if and only if synergies are sufficiently large—namely, above $\hat{s} = (a-c)/3c$. The second is a dynamic policy that approves a merger if and only if long-run average consumer surplus increases. Letting $\theta(\beta_0,\beta_1) = (\beta_0 - \beta_1)/(1-\beta_1)$ denote the probability of successful innovation given initial beliefs β_0 and the termination of R&D at beliefs $\beta_1 < \beta_0$, this dynamic review policy finds approves a merger if and only if

$$\theta(\beta_0,\beta^M)CS_M^+ + (1-\theta(\beta_0,\beta^M))CS_M(0) > \theta(\beta_0,\beta_1^*)CS^+ + (1-\theta(\beta_0,\beta_1^*))CS(0,0).$$

Upon rearrangement, we observe that dynamic review approves a merger if and only if the sum of the merger's combined static and dynamic effects on consumer surplus are positive:

$$\underbrace{CS_M(0) - CS(0,0)}_{\text{Static Effect}} + \underbrace{\theta(\beta_0,\beta^M)[CS_M^+ - CS_M(0)] - \theta(\beta_0,\beta_1^*)[CS^+ - CS(0,0)]}_{\text{Dynamic Effect}} > 0.$$

A key policy question is how static and dynamic review compare in terms of enforcement. Using Corollary 2, we note that merger's dynamic effect is decreasing—that is, the merger becomes less desirable from a dynamic perspective. Therefore, to obtain a comparison between static and dynamic merger review, it suffices to determine the sign of the dynamic effect at $s = \hat{s}$. If the dynamic effect is negative at $s = \hat{s}$, then static review is overly permissive in the sense that it approves mergers that appear optimal from a static perspective, but nevertheless reduce long-run average consumer surplus. Likewise, static review is overly restrictive whenever the dynamic effect is positive at $s = \hat{s}$ because mergers that increase long-run average consumer surplus are blocked due to their negative short-run effects.

Proposition 19 (Mergers and Welfare)**.**

1. *Suppose that the merger involves consolidation. Then static merger review is overly permissive compared to dynamic merger review for all $\phi \geq 0$.*
2. *Suppose that the merger involves coordination. Then static merger review is overly permissive (restrictive) compared to dynamic merger review whenever ϕ is sufficiently small (large).*

Proposition 19 shows that the comparison between static and dynamic merger review critically depends two factors: (i) whether there is coordination or consolidation, and (ii) the endogenous opportunity cost of R&D, as measured by ϕ. In the case of a merger with consolidation, we know from Proposition 18 that Firm M, absent synergies, has lower incentives to innovate compared to duopoly. By Corollary 2, an increase in s only makes Firm M's innovation incentives lower. Therefore, the merger's effect on long-run average consumer surplus is always lower than the merger's static effect, and static review will approves more mergers than is dynamically optimal.

A comparison between static and dynamic merger effects is more difficult to establish when firms coordinate after merging. The reason is that the market structure and innovation incentive effects can have opposite sign. In particular, even if there is a static reduction in consumer surplus due to low synergies, there may be dynamic gains from increased innovation incentives. The question is whether innovation incentives under coordination remain higher than duopoly at $s = \hat{s}$. As the proof shows, the answer to this question depends on how beneficial coordination is to Firm M's incentives to innovate which. If ϕ is relatively small, the opportunity cost of R&D in duopoly is close to zero, and a merger with coordination produces, at best, only a small reduction in the opportunity cost of R&D. As a result, Firm M's incentive to innovate will be lower than duopoly incentives at $s = \hat{s}$ in this case. But if ϕ is large, then coordination produces a large reduction in the opportunity cost of R&D and, hence, a large increase in innovation incentives. In this case, Firm M's incentive to perform R&D remains sufficiently above duopoly incentives for all $s < \hat{s}$ to guarantee a positive dynamic effect over this range of synergies. As a result, static merger review permits too few mergers because it underestimates the value of coordination for lowering the opportunity cost of R&D.

Proposition 19 highlights an important nuance regarding the ability of competition authorities to make presumptions about consumer harm through dynamic innovation effects. In its decision in the *Dow/DuPont* case, the European Commission articulated an innovation theory of harm that concluded that mergers generally stifle innovation and, therefore, can be presumed to be anti-competitive absent large static benefits (see Denicolo and Polo, 2018, for related discussion). My results, in particular Propositions 18 and 19, provide conditions under which such a presumption of harm can reasonably be applied. In particular, if a merger is expected to produce significant consolidation, then the effects of enhanced market power are likely to reduce R&D incentives. In such cases, it would be appropriate for competition authorities to apply stricter merger standards or require larger remedies to account for the merger's negative dynamic effects.

If consolidation is unlikely to occur, however, then a presumption of harm to innovation incentives and welfare may lead to overly restrictive merger enforcement that that blocks dynamically pro-competitive mergers. In some cases, such as a "killer acquisition" (see Cunningham et al., 2021), my results suggests that a presumption of harm to both innovation and welfare is appropriate. In others, such as common ownership (see Anton et al., 2021), the appropriateness of a presumption of harm depends on industry and firm specifics. Therefore, to maximize both innovation and consumer welfare in competitive industries, is required that competition authorities take a case-by-case view in assessing a merger's anticipated effects on welfare, as the comparison between static and dynamic effects crucially depend on the features of both market structure and the proposed merger.

3.7 Concluding Remarks

This dissertation chapter developed a tractable framework to analyze the role of endogenous product market considerations on R&D incentives and competition. I have shown that both the positive and normative implications of R&D investment vary considerably depending on the details about market structure and, in particular, the value of alternative investment. Finally, I have illustrated the applicability of my framework to current policy issues, such as the design of innovation subsidies and the evaluation of mergers.

The basic framework I develop is flexible and can be extended in several dimensions to analyze more specific questions and industry contexts. For example, one may wish to assume more structure on terminal payoffs following successful innovation to examine other important policy questions. Examples could include the length or breadth of patents (Gilbert and Shapiro, 1990), ex-ante or ex-post licensing (Katz and Shapiro, 1985), and rewards to sequential innovation Hopenhayn et al. (2006). For any of these extensions, my results will inform the structure of MPE experimentation and its overall social efficiency. However, one may be able to obtain greater tractability with more structure and thus obtain more specific, policy-relevant comparative statics.

Other extensions would be more substantial but could interesting results. First, by focusing on the case of binary effort, my framework abstracts from the scale of investment. Introducing a choice of scale would greatly complicate my analysis since flow payoffs in MPE would then directly depend on current beliefs, resulting in highly nonlinear value functions. Nevertheless, such a modification would allow one to ask important questions related to firm size (i.e., scale choice) and the nature of investment (Cohen and Klepper, 1996). Second, to highlight the fact the role of market structure in generating asymmetric R&D incentives, my analysis assumes that firms are ex-ante symmetric. Interacting this endogenous asymmetry with other asymmetries, such as initial quality or cost advantage, may uncover useful insights related to innovation and market dominance (Athey and Schmutzler, 2001). Such asymmetry would undoubtedly complicate the analysis of value functions. However, it may be possible to retain tractability with appropriate simplifications to the base model.

3.8 Appendix: Proofs and Derivations

Proofs of Proposition 10 and Lemma 7. Given in text. □

Proof of Proposition 11. First note that $\overline{V}_{(0,1)}(\beta)$ and $(1-\beta)\Phi(\beta)^{\frac{\rho}{\lambda}}$ are decreasing in β. This implies that $V_{(0,1)}(\beta)$ is decreasing over $(\beta_1^*, 1)$ whenever $\mathbb{C}_{(0,1)}$ is positive. Hence, an innovator disadvantage exists only if $\mathbb{C}_{(0,1)}$ is negative. But whenever $\mathbb{C}_{(0,1)}$ is negative, we know that $V_{(0,1)}(\beta)$ is strictly concave. Therefore, if an innovator disadvantage exists, it will exist over an interval $(\beta_1^*, \hat{\beta})$ for some $\hat{\beta} > \beta_1^*$. By strict concavity, a necessary and sufficient condition for $\hat{\beta} > \beta_1^*$ is that $V'_{(0,1)}(\beta_1^*+) > 0$. Using the Bellman equation that defines $V_{(0,1)}(\beta)$, take the limit of $V'_{(0,1)}(\beta)$ as $\beta \to \beta_1^*$ from above to obtain

$$\Pi(0,0) = \Pi(0,1) + \beta_1^*\lambda[V^- - V_{(0,0)}] - \lambda\beta_1^*(1-\beta_1^*)V'_{(0,1)}(\beta_1^*+).$$

Observe from this equation that $V'_{(0,1)}(\beta_1^*+)$ is positive iff $\Pi(0,0) < \Pi(0,1) + \beta_1^*\lambda[V^- - V_{(0,0)}]$ which, in turn, simplifies to give our desired condition.

To complete the proof, we must establish that if an innovator disadvantage exists, then it exists over an interval $(\beta_1^*, \hat{\beta})$. To do this, it suffices to verify that $V_{(0,1)}(\beta)$ is either strictly convex, strictly concave, or linear, since $V_{(1,0)}(\beta)$ is strictly convex.

The general solution to (39) is $V_{(0,1)}(\beta) = \overline{V}_{(0,1)}(\beta) + \mathbb{C}_{(0,1)}(1-\beta)\Phi(\beta)^{\frac{\rho}{\lambda}}$, where

$$\overline{V}_{(0,1)}(\beta) = (1-\beta)\left(\frac{\Pi(0,1)}{\rho}\right) + \beta\left(\frac{\Pi(0,1)+\lambda V^-}{\rho+\lambda}\right),$$

and $\mathbb{C}_{(0,1)}$ is a constant of integration. As $\beta \to \beta_1^*$, value matching implies that $V_{(0,1)}(\beta_1^*) = V_{(0,0)}$. Hence, the exact solution for $\mathbb{C}_{(0,1)}$ is

$$\mathbb{C}_{(0,1)} = \frac{1}{(1-\beta_1^*)\Phi(\beta_1^*)^{\frac{\rho}{\lambda}}}\left(V_{(0,0)} - \overline{V}_{(0,1)}(\beta_1^*)\right).$$

Insert this into the general solution for $V_{(0,1)}(\beta)$ to obtain the following exact solution:

$$V_{(0,1)}(\beta) = \overline{V}_{(0,1)}(\beta) + \frac{(1-\beta)\Phi(\beta)^{\frac{\rho}{\lambda}}}{(1-\beta_1^*)\Phi(\beta_1^*)^{\frac{\rho}{\lambda}}}\left(V_{(0,0)} - \overline{V}_{(0,1)}(\beta_1^*)\right).$$

The function $\overline{V}_{(0,1)}(\beta)$ is linear in β, and $(1-\beta)\Phi(\beta)^{\rho/\lambda}$ is strictly convex in β. Thus, $V_{(0,1)}(\beta)$ is either strictly concave, strictly convex, or linear in β. In either case, $V_{(0,1)}(\beta)$ and $V_{(1,0)}(\beta)$ are strictly single crossing over $(\beta_1^*, 1)$, which establishes that if an innovator disadvantage exists, it must exist over an interval $(\beta_1^*, \hat{\beta})$ for some $\hat{\beta} \in (\beta_1^*, 1]$. □

Proof of Propositions 12 and 13. The proof of this result can be broken into three steps.

Lemma 8. *Simultaneous exit must occur at* β_1^* *must occur in any MPE whenever* $RE > CE$.

Proof. Suppose that $RE > CE$. We know from our previous discussions—namely, the comparison of β_2^* and β_1^*—that sequential exit cannot occur in MPE. The question is whether simultaneous exit can occur below β_1^*. Suppose that simultaneous exit occurs at the belief $\tilde{\beta}$ in a particular MPE. Clearly $\tilde{\beta} \leq \beta_1^*$ must hold; otherwise, one firm would have a strict incentive to deviate. Let $V_{\tilde{\beta}}(\beta)$ denote the expected value of Firm i in such an MPE. Since there is option value to staying in the race, we know that $V_{\tilde{\beta}}(\beta)$ is strictly convex over $(\tilde{\beta}, 1)$. If it were true that $\tilde{\beta} < \beta_1^*$, then exit by Firm i would immediately induce exit by Firm j for all $\beta \in (\tilde{\beta}, \beta_1^*)$. Therefore, $V_{(1,1)}(\beta) \geq V_{(0,0)}$ must hold for all $\beta \in (\tilde{\beta}, \beta_1^*)$ in this MPE. This inequality, along with strict convexity, implies that $V'_{\tilde{\beta}}(\tilde{\beta}+) \geq 0$ must hold in MPE.

To determine conditions under which $V'_{\tilde{\beta}}(\tilde{\beta}+) \geq 0$ is possible in MPE, take the limit of Equation (41) as beliefs approach $\tilde{\beta}+$ to obtain

$$\Pi(0,0) = \Pi(1,1) - f + \tilde{\beta}\lambda[V^+ + V^- - 2V_{(0,0)}] - 2\lambda\tilde{\beta}(1-\tilde{\beta})V'_{\tilde{\beta}}(\tilde{\beta}+).$$

From this equation it follows that $V'_{\tilde{\beta}}(\tilde{\beta}+) \geq 0$ holds iff

$$\Pi(0,0) \leq \Pi(1,1) - f + \tilde{\beta}\lambda\left[V^+ + V^+ - 2V_{(0,0)}\right] \iff \tilde{\beta} \geq \frac{\rho}{\lambda}\left(\frac{\Pi(0,0) - \Pi(1,1) + f}{\Pi^+ + \Pi^- - 2\Pi(0,0)}\right).$$

Thus, a necessary condition for $\tilde{\beta} < \beta_1^*$ is

$$\beta_1^* > \frac{\rho}{\lambda}\left(\frac{f + \Pi(0,0) - \Pi(1,1)}{\Pi^+ + \Pi^- - 2\Pi(0,0)}\right) \iff \Pi(1,1) - \Pi(1,0) > \left(\frac{\Pi(0,0) - \Pi^-}{\Pi^+ - \Pi(0,0)}\right)(f + \mathscr{C}(0)).$$

We know by submodularity that $\Pi(0,1) - \Pi(0,0) \geq \Pi(1,1) - \Pi(1,0)$. Thus, the above necessary condition holds only if there exists an innovator disadvantage. But we have already shown that a necessary condition for an innovator disadvantage to exist is $CE > RE$. Thus, the only possible value MPE outcome when $RE > CE$ is that both firms simultaneously exit at $\tilde{\beta} = \beta_1^*$. □

Lemma 9. *Sequential exit must occur at $\hat{\beta}$ in any MPE whenever $CE > RE$ and $\hat{\beta} \geq \beta_2^*$.*

Proof. Suppose that $CE > RE$ and that $\hat{\beta} \geq \beta_2^*$. We know from in-text discussions that sequential exit must occur in any MPE. For all $\beta > \hat{\beta}$, it is easy to see that both firms stay in the race with probability one since $\beta > \beta_2^*$ implies that $e_i = 1$ is a best response to e_j over $(\hat{\beta}, 1)$. So the question becomes if sequential exit can possibly occur below $\hat{\beta}$.

Suppose that Firm j exits at the belief $\beta^j < \hat{\beta}$ if Firm i were still in the race. Given this exit strategy, Firm i has a strict incentive to exit the race at $\beta^j + \epsilon$ for sufficiently small $\epsilon > 0$ instead of stay in the race because $V^*_{(0,1)}(\beta^j) > V^*_{(1,0)}(\beta^j)$. But if we swap the roles of Firm i and Firm j in this best-response calculation, we see that Firm j has the same incentive given that Firm i exits at $\beta^j + \epsilon$. As in Fudenberg and Tirole (1985), we can repeatedly apply this preemption logic to conclude that sequential exit must occur $\hat{\beta}$ in MPE. □

Lemma 10. *Suppose that $CE > RE$ and that $\beta_2^* > \hat{\beta}$. Then the timing of sequential exit is determined via a war of attrition as soon as beliefs reach β_2^*. In this war of attrition, either (a) one firm exits immediately or (b) each firm randomly exit with intensity*

$$\mu(\beta) = \frac{\Pi(0,0) + CE - \widetilde{RE}(\beta) - \rho V_{(0,1)}(\beta)}{V_{(1,0)}(\beta) - V_{(0,1)}(\beta)}, \text{ where } \widetilde{RE}(\beta) = \beta\lambda(V^+ - V^-) - f.$$

Proof. Suppose that $CE > RE$ and that $\beta_2^* > \hat{\beta}$. We know from in-text discussions that sequential exit must occur in any MPE. For all $\beta > \beta_2^*$, it is easy to see that both firms stay in the race with probability one since $\beta > \beta_2^*$ implies that $e_i = 1$ is a best response to e_j over $(\beta_2^*, 1)$. So the question becomes if sequential exit can possibly occur below β_2^*.

For all $\beta \in (\hat{\beta}, \beta_2^*)$, the subgame in which both firms are in the race is strategically equivalent to a continuous-time war of attrition, as defined in Hendricks et al. (1988). Adapting their results to the present setting, one concludes that three MPE outcomes exist. Two outcomes feature sequential exit at β_2^* with probability one, and one outcome features symmetric mixing over $(\hat{\beta}, \beta_2^*)$. Let $\mu(\beta)$ denote the intensity at which each firm exits in this latter MPE. To determine its formula, write out the Bellman equation defining the value of Firm i whenever randomization occurs:

$$V_i(\beta) = \Pi(1,1) - f + \beta\lambda\left[V^+ + V^+ - 2\tilde{V}_i(\beta) - 2(1-\beta)V_i'(\beta)\right] + \mu(\beta)\left[V_{(1,0)}(\beta) - V_i(\beta)\right].$$

By definition, Firm i must be indifferent between $e_i = 1$ and $e_i = 0$ for all $\beta \in (\hat{\beta}, \beta_{(1,1)}^*)$. Therefore, the value of Firm i satisfies $V_i(\beta) = V_{(0,1)}(\beta)$ and $V_i'(\beta) = V_{(0,1)}'(\beta)$ for all $\beta \in (\hat{\beta}, \beta_2^*)$. Inserting these boundary conditions into the Bellman equation yields

$$\mu(\beta) = \frac{\Pi(0,0) + CE - \widetilde{RE}(\beta) - \rho V_{(0,1)}(\beta)}{V_{(1,0)}(\beta) - V_{(0,1)}(\beta)}, \text{ where } \widetilde{RE}(\beta) = \beta\lambda(V^+ - V^-) - f.$$

To finish the proof, note that the denominator of $\mu(\beta)$ vanishes as β approaches $\hat{\beta}$, but the numerator is bounded above zero. Therefore, $\mu(\beta)$ tends to infinity as beliefs approach $\hat{\beta}$, which implies that sequential exit must occur with probability one before $\hat{\beta}$. $\square$

Proof of Proposition 14. Whenever an innovator disadvantage exists, the belief $\hat{\beta}$ sup$\{\beta \in (0,1) : V_{(0,1)}(\hat{\beta}) \geq V_{(1,0)}(\hat{\beta})$ is greater than β_1^*. If interior, $\hat{\beta}$ can be re-written as

$$\overline{V}_{(0,1)}(\hat{\beta}) - \overline{V}_{(1,0)}(\hat{\beta}) = \frac{(1-\hat{\beta})\Phi(\hat{\beta})^{\frac{\rho}{\lambda}}}{(1-\beta_1^*)\Phi(\beta_1^*)^{\frac{\rho}{\lambda}}} \left(\overline{V}_{(0,1)}(\beta_1^*) - \overline{V}_{(1,0)}(\beta_1^*)\right).$$

The function $(1-\beta)\Phi(\beta)^{\frac{\rho}{\lambda}}$ is strictly decreasing in β since

$$\frac{d}{d\beta}(1-\beta)\Phi(\beta)^{\frac{\rho}{\lambda}} = -\left(1 + \frac{\rho}{\beta\lambda}\right)\Phi(\beta)^{\frac{\rho}{\lambda}} < 0.$$

Thus, the ratio between $(1-\hat{\beta})\Phi(\hat{\beta})^{\frac{\rho}{\lambda}}$ and $(1-\beta_1^*)\Phi(\beta_1^*)^{\frac{\rho}{\lambda}}$ is less than one. As ρ increases, we know that β_1^* eventually crosses 1. For the right-hand side of the first equation to be defined as $\beta_1^* \to 1$, it must be true that $\hat{\beta} = 1$ already holds. Let $\bar{\rho}$ denote the lowest discount rate at which $\hat{\beta} = 1$ holds. The previous argument guarantees that $\beta_1^* < 1$ holds for all $\rho > \bar{\rho}$. Over this range of discount rates, we necessarily have $\hat{\beta} \geq \beta_2^*$. Thus, preemption is the unique MPE outcome for all $\rho > \bar{\rho}$ such that $\beta_1^* < \beta_0$, as desired. $\square$

Proof of Proposition 15. Provided that sequential exit occurs in the planner's optimal experimentation policy, the construction of β_1^S and β_2^S are analogous to β_1^* and β_2^*. Therefore, it suffices to prove that optimal experimentation always prescribes sequential exit.

Given $\beta \in (0,1)$ and $(e_i, e_j) \in \{0,1\}$, the HJB equation defining the social planner's value, denoted $V_S(\beta)$, is

$$\rho V_S(\beta) = TS(e_i, e_j) - f(e_i + e_j) + \beta\lambda(e_i + e_j)\left[V_S^+ - V_S(\beta) - (1-\beta)V_S'(\beta)\right],$$

where $V_S^+ \equiv TS^+/\rho$ denotes the social value of successful innovation.

To derive the indifference condition for optimal stopping given $e_j = 1$, note that the HJB equation defining $V_S(\beta)$ can be re-written as $\rho V_S(\beta) = TS(0, e_j) + e_j(B(\beta, V_S) - f) + e_i(B(\beta, V_S) - f - C_S(e_j))$, where $B(\beta, V_S) \equiv \beta\lambda[V_S^+ - V_S(\beta) - (1-\beta)V_S'(\beta)]$, and $C_S(e_j) \equiv TS(0, e_j) - TS(1, e_j)$. The planner is indifferent between the actions $e_i \in \{0, 1\}$ if and only if $B(\beta, V_S) = f + C_S(e_j)$, which is equivalent to the condition $\rho V_S(\beta) = TS(0, e_j) + e_j C_S(e_j)$.

By direct inspection, the planner's joint R&D incentive is lower than its individual R&D incentive if and only if $TS(0,1) + C_S(1) > TS(0,0)$, which simplifies to give $2TS(0,1) > TS(0,0) + TS(1,1)$. Consequently, the optimal R&D policy features sequential whenever $TS(e_i, e_j)$ is strictly submodular. We know that both $\Pi(e_i, e_j)$ is strictly submodular whenever $\phi > 0$. Therefore, all we must check is whether $\partial^2 \frac{1}{2}(Q(e_i, e_j))^2 / \partial e_i \partial e_j \leq 0$ holds. Computing this derivative, we obtain $\frac{\partial^2}{\partial e_i \partial e_j}\frac{1}{2}Q(e_i, e_j)^2 = -\frac{\phi^2}{9} \leq 0$. Thus, sequential exit occurs iff $\phi > 0$. ☐

Proof of Propositions 16 and 17. To determine whether the total amount of R&D performed in equilibrium is, on average, higher or lower than the social optimum, it suffices to compare β_1^S to β_1^*. Using the decomposition $\beta_1^S = \beta_f^S + \beta_\phi^S$, where β_f^S and β_ϕ^S denote the explicit and implicit components of β_1^S, we can compare the implicit components as follows:

$$\beta^S_\phi = \beta^c \left(\frac{8(a - c_A - \phi) - 3\phi}{8(a - c_B) + 3(c_A - c_B)} \right) < \beta^*_\phi.$$

Furthermore, through direct inspection, we have $TS^+ - TS(0,0) > \Pi^L - \Pi(0,0)$ for all relevant parameter values. Therefore, we have $\beta^S_f < \beta^*_f$ as well, which allows us to unambiguously conclude that $\beta^S_1 < \beta^*_1$ holds for all relevant parameter values.

To determine the efficiency of equilibrium R&D intensity for small discount rates, we make use of the following limit results, each of which can be easily verified:

1. $\lim_{\rho \to 0} \beta^*_1 = 0$.
2. $\lim_{\rho \to 0} \rho V^*_{(0,1)}(\beta) = (1 - \beta)\Pi(0,0) + \beta \Pi^-$ (i.e. discounted value approaches the long-run average).
3. $\lim_{\rho \to 0} \rho V^*_{(1,0)}(\beta) = (1 - \beta)\Pi(0,0) + \beta \Pi^+$ (i.e. discounted value approaches the long-run average).

For all $\rho > 0$, we have $\rho V^*_{(0,1)}(\beta^*_2) = \Pi(0,1) + C(1) - \beta^*_2 \lambda (V^L - V^F)$, by definition. In the limit as $\rho \to 0$, the left-hand side of this equation remains finite (by the above limit result). Therefore, the right-hand side must also remain finite, which implies that $\lim_{\rho \to 0} \beta^*_2 = 0$ because $V^L - V^F$ increases without bound as discounting becomes small. Likewise, $\hat{\beta}$ is defined by the equation $\rho V^*_{(1,0)}(\hat{\beta}) = \rho V^*_{(0,1)}(\hat{\beta})$. As $\rho \to 0$, this inequality must continue to hold. However, by our above limit result, the only way this is possible is if $\lim_{\rho \to 0} \hat{\beta} = 0$. From this, we conclude that $\max\{\beta^*_2, \hat{\beta}\}$ approaches zero as discounting vanishes.

Now consider the limit $\lim_{\rho \to 0} \beta^S_2$. By definition, β^S_2 is defined by the indifference condition $\rho V_S(\beta_2^S) = TS(0,1) + C_S(0,1) > 0$. As $\rho \to 0$, the planner's discounted value (for analogous reasons to our second limit result) converges to $(1-\beta)TS(0,0) + \beta TS^+$. Consequently, we have

$$\lim_{\rho \to 0} \beta_2^S = \frac{TS(0,1) - TS(0,0) + C_S(0,1)}{TS^+ - TS(0,0)} > 0.$$

Therefore, as $\rho \to 0$, we will eventually have $\beta_2^S > \max\{\beta_2^*, \hat{\beta}\}$. So the intensity of equilibrium R&D, in any Markov-perfect equilibrium, will be, on average, higher than is socially desired. □

Proof of Corollary 2. First, suppose that the merger does not involve consolidation. Given $\beta_0 > \beta_{Coord}^M$, the total amount of R&D Firm M performs before exiting the patent race is

$$N_{Coord}^M = \int_{\beta_{Coord}^M}^{\beta_0} \frac{d\beta}{\lambda\beta(1-\beta)} = \frac{1}{\lambda}\left[\log\Phi(\beta_{Coord}^M) - \log\Phi(\beta_0)\right].$$

This expression is strictly decreasing in β_{Coord}^M. Therefore, it sufficies to check that β_{Coord}^M is increasing in s. Through direct inspection, we see that this is true since

$$\beta_{Coord}^M = \frac{\rho}{\lambda}\left(\frac{4bf}{(1-s)(c-c^+)(2a-(1-s)(c+c^+))}\right).$$

Now suppose that the merger involves consolidation. Given $\beta_0 > \beta_{Cons}^M$, the total amount of R&D Firm M performs before exiting the patent race is

$$N^M_{Cons} = \int_{\beta^M_{Cons}}^{\beta_0} \frac{d\beta}{\lambda\beta(1-\beta)} = \frac{1}{\lambda}\left[\log\Phi(\beta^M_{Cons}) - \log\Phi(\beta_0)\right].$$

As before, this expression is strictly decreasing in β^M_{Cons}. Therefore, it sufficies to check whether β^M_{Cons} is increasing or decreasing in s. Upon simplification, we obtain

$$\beta^M_{Cons} = \frac{\rho}{\lambda}\left(\frac{4bf}{(1-s)(c-c^+)(2a-(1-s)(c+c^+))}\right) + \frac{\rho}{\lambda}\left(\frac{\phi}{c-c^+}\right)\left(\frac{2a-(1-s)(2c+\phi)}{2a-(1-s)(c+c^+)}\right).$$

The first term in β^M_{Cons} is increasing in s since the numerator is constant in s and the demonator is decreasing in s. Likewise, the second term in β^M_{Cons} is increasing in s since $c+\phi > c^+$. Indeed, the derivative of the second term with respect to s satisfies

$$\frac{\rho}{\lambda}\left(\frac{2a\phi(c+\phi-c^+)}{(c-c^+)(2a-(1-s)(c+c^+))^2}\right) > 0 \iff c+\phi > c^+,$$

which completes the proof. □

Proof of Proposition 19. A merger is optimal from the standpoint of long-run average consumer surplus if and only if $CS_M(0) + \theta(\beta_0, \beta^M)\big(CS^+_M - CS_M(0)\big)$ is greater than $CS(0,0) + \theta(\beta_0, \beta^*_1)(CS^+ - CS(0,0))$, where $\theta(\beta_0, \beta_1) = (\beta_0 - \beta_1)/(1-\beta_1)$ denotes the ex-ante probability of successful innovation given initial beliefs β_0 and that a non-zero amount of R&D is performed until beliefs cross β_1. Through rearrangement, this inequality can be re-written a requirement that the *Static Effect* plus the *Dynamic Market Structure*

Effect plus the *R&D Incentive Effect* is non-negative, where

1. *Static Effect* $= CS_M(0) - CS(0,0)$
2. *Dynamic Market Structure Effect* $= \theta(\beta_0, \beta_1^*)\left[(CS_M^+ - CS_M(0)) - (CS^+ - CS(0,0))\right]$
3. *R&D Incentive Effect* $= \left[\theta(\beta_0, \beta^M) - \theta(\beta_0, \beta_1^*)\right](CS_M^+ - CS_M(0))$

Both the dynamic market structure and R&D incentive effects are decreasing in s. Thus, a necessary and sufficient condition for static merger review to be overly permissive is that $CS_M(0) = CS(0,0)$ implies that each dynamic effect is negative.

Direct calculation shows that $CSM(0) > CS(0,0)$ if and only if $a < 4c$ and $s > \frac{a-c}{3c}$, as previously noted in the text. Suppose that $CS_M(0) = CS(0,0)$. Then the dynamic market structure effect is equal to $\theta(\beta_0, \beta_1^*)\left[CS_M^+ - CS^+\right]$. Direct calculation shows that $CS_M^+ > CS^+$ if and only if $s > \frac{a-2c+c^+}{3c^+}$. Finally, direct inspection shows that $\frac{a-2c+c^+}{3c^+} > \frac{a-c}{3c}$ always holds. Thus, the dynamic market structure effect is negative if $CS_M(0) = CS(0,0)$. Finally, we can use our previous results to conclude that the R&D incentive effect is always negative since $\theta(\beta_0, \beta_1)$ is decreasing in β_1 and $\beta^M > \beta_1^*$ holds for all $s \in [0,1]$. □

Chapter 4

Persistence and the Value of Leadership

4.1 Introduction

When an R&D project is initiated, it is understood that success may be a long time in coming…if it comes at all. In this context a key question for management is whether to continue to support the project or to reallocate resources to a different project, thereby starving or terminating the existing project.[1] Such resource decisions are typically made during periodic reviews of the attractiveness of the current project and other potential R&D prospects.[2] In the drug R&D process, for example, projects proceed through sequential reviews corresponding to pre-clinical and clinical phases of each project. Termination is a common outcome.[3]

The possibility of termination also affects effort exerted on the project which, in turn,

---

[1] One practitioner put the problem as follows: "[T]erminating a research and development project is one of the most agonizing decisions that confronts R&D administrators because those working on a problem are invariably sure that success is just around the corner. However, like drilling oil wells, it is not practical to keep on drilling forever." (Buell 1967) Another underscored the importance of the decision: [In managing pharmaceutical project portfolios] "[t]he most important, and at the same time the most difficult, decision is dealing with project termination." (Jekunen 2014)

[2] Another example of periodic review is the "stage-gate" methodology which many firms use as part of their decision-making process (Cooper, Edgett, and Kleinschmidt 1998). This methodology is designed around a series of project phases and "decision gates" that allow a project to proceed to the next phase. An analogous sequential process is followed by venture capitalists when they review the projects that they fund.

[3] The frequency of termination varies widely. While termination is quite common in R&D pharmaceutical projects (Jekunen 2014, Krieger 2021) it is less common for less risky projects. For example, in a general survey of Japanese companies it was found that about 60% had terminated or suspended one project in the three years preceding the authors' survey (Haneda and Ono 2022) while in a German firm survey, 1.37 projects were terminated while 8.38 were initiated over a three-year period (Andries and Hünermund 2020).

affects the termination decision itself. If projects are terminated quickly, subordinates may become skeptical about a leader's willingness to commit to long-term projects. In their study of corporate resource allocation decisions, for example, Bower and Gilbert 2005 (p.11) found "...there can be paralysis, a constant debate where no strategy is actually implemented with commitment and continuity, and where operating managers cynically wait for the next 'flavor of the month' and deflect one strategic initiative after another." This skepticism has attendant consequences for the subordinates' own commitments to such projects and is a major concern of leaders.[4] A leader's strategy for persistence and termination, then, not only affects a firm's R&D direction, it is a policy instrument for gaining commitment and motivating effort of teams doing the work of the organization.

This dissertation chapter explores the value of the persistence instrument and its impact on subordinate effort. We build and analyze an analytical model of the R&D project process which focuses on the persistence decision. Our equilibrium analysis combines an optimal search problem (When should a leader abandon a current project in favor of a new project?) with an agency problem (Can persistence induce subordinates to exert greater effort on the selected project?).

We model a leader's decision whether to terminate an existing R&D project in favor of a new R&D project with an uncertain value. This scenario corresponds to a setting in which a firm is carrying only a single project in the early development stage but has other

---

[4] Subordinate commitment is commonly listed as a factor management considers when deciding whether to persist with a project. Other criteria include financial factors such as anticipated returns and risk that depend on assessments of technical feasibility, costs, market potential and acceptance as well as factors such as strategic fit and support of various levels of management. (Balachandra and Raelin 1980; Cooper, Edgett, Kleinschmidt 1998). The relative importance of any given factor depends on the stage at which termination is contemplated.

even earlier stage research projects to which it could switch. To capture the search and agency factors, we build a discrete-time, infinite-horizon model of the project search and development environment in which project success depends on effort. In each period, the leader has control over whether the organization starts a new project or persists with the existing project. Once a project is initiated the leader learns the value of a successful project. Each period consists of three stages. First, the subordinate (henceforth, manager) decides whether to devote effort to develop the existing project. Success is uncertain, but only is possible with manager effort and effort is not observable to the leader. Then, the outcome of effort is realized. If development was successful, a new project with an uncertain value is initiated (drawn from a distribution of projects in which projects differ only by their ultimate commercial value given success). If development was unsuccessful, the leader chooses whether to persist with the existing project. If the existing project is terminated, then a new project is initiated.

We seek to understand the potential for persistence by the leader to influence incentives for followers to provide (costly) effort. This focus leads us to abstract from compensation and to suppose that followers are rewarded by sharing in the value of a successful project. The model highlights the essential feature of the interaction between persistence and effort in the model: if a leader is expected to persist with a project, then the only way for the manager to move on to a new project is to generate a success via effort since shirking ensures a failure today and then persistence by the leader with this project for the next period. The simplicity of this core model also allows us to explore interesting variations built on the same structure.

We begin with an analysis in which both the leader and the manager know the value of the current project being developed, but not the value of the future projects. This complete information setting highlights the importance of persistence and the attractiveness of the underlying distribution of future projects on equilibrium outcomes. The lure of future project payoffs affects the bottom end of the range of projects on which the leader will persist. If the value distribution of future projects is sufficiently favorable relative to the costs of effort, that is, if the firm faces a “bright future,” the leader will terminate a range of projects that would pass a static cost-benefit benchmark. This aggressive termination strategy contrasts with the highly persistent strategy which obtains if the distribution of future projects is unfavorable (a “dim future”). In the latter case, persistence induces manager effort on lower-value projects, including a range that would fail the static cost-benefit benchmark, thereby making such projects more attractive to the leader.[5] Absent persistence, the manager’s choice of effort depends on whether the payoff to a project justifies effort on that project. But with persistence, the manager can be induced by the leader to exert effort because success with the previous project is the only way the manager can gain access to attractive new projects.

---

[5] As an example, consider how the pharmaceutical firm Merck ’s future, as reflected in the thickness of the early-stage R&D drug pipeline, affected their project selection criteria. “[Merck’s] history of policy shifts suggests at a minimum that Merck’s financial hurdles for advancing projects fluctuated over the past 20 years—shifting from high thresholds selecting only to develop potential blockbuster drugs, to lower thresholds that included Merck’s willingness to commercialize its less valuable projects and to acquire other firms’ discoveries.” (Chan et. al. 2007). More generally, but also for pharmaceutical firms, Andries and Hunermund (2020) finds resource-abundant firms start more new projects and abandon more projects than resource-constrained firms. Their preferred explanation is that availability of resources leads to over-optimism which in our terms can be interpreted as a leader thinking the firm has a “brighter future.”

Next, we provide a complete characterization of equilibria in the setting in which only the leader knows the expected commercial value of the project, presumably because of greater knowledge about the organization's commercialization activities, other ongoing projects, competitive intelligence, and so on.[6] How might this private information affect the impact of the persistence instrument on subordinate effort and project outcomes?

Private information is introduced to the core model by assuming the leader learns the value of the existing project whereas, ex ante, the manager only knows the underlying distribution of project values. The persistence choice, then, may provide a signal to the manager about the leader's private information which leads to an updating of the manager's beliefs about the existing project's value. Equilibrium outcomes in this asymmetric information analysis depend on the mean and shape of the underlying ex ante distribution of project values. A particularly interesting "wait and see" equilibrium occurs where the manager withholds effort when a project is first introduced, then works when the leader persists with the (necessarily) unsuccessful project. This equilibrium is driven by two factors. First, the manager's ex ante belief regarding the project value is sufficiently poor that initial-period effort has low or negative expected value. Second, the leader's choice to persist with the existing project credibly signals to the manager that the expected value of the project is sufficiently high to now justify effort because the manager correctly anticipates

---

[6] These aspects of project valuation are especially emphasized by studies of project portfolio choices. For example, in discussing R&D strategy more generally, Brennan et. al. (2020) argue "[d]eveloping a strategy for the R&D organization entails some unique challenges that other functions do not face. For one, scientists and engineers have to weigh considerations beyond their core expertise, such as customer, market, and economic factors." Others (e.g., Helfat and Eisenhardt 2004, Levinthal and Wu 2010) note the importance of freeing firm resources for other opportunities which presumably would be better understood by higher-level managers than the manager of the team specifically tasked with one or two projects.

that the leader only persists with higher-value projects. This updated belief justifies manager effort which, in equilibrium, also justifies the leader's choice to persist on higher-value projects. More subtly, because the persistence signal only applies when the initial period outcome is a failure, there can be a gain to the manager of waiting and ensuring failure. To see this, compare the situation when the manager exerts effort in the initial period to that when the manager waits. When the result of effort is failure (despite effort in the first case and because of no effort in the second), then both situations offer the same future expected payoff streams. If the leader persists, then the payoff stream begins with a higher-value project on average. If not, then the payoff stream begins with a new project. In contrast, when effort results in success, the manager gains the expected value of the initial project (minus the cost of effort) and a discounted payoff stream beginning with a new project. But because the expected payoffs from a new project stream are less than from a stream beginning with a higher-expected-value project, waiting may dominate exerting effort. This effect is sufficiently strong that the equilibrium threshold for persistence can exceed the static benchmark.

Incomplete information adds cost relative to the complete information benchmark because for sufficiently unfavorable underlying distributions of project value, the leader cannot induce effort in the first period of a new project. The wait and see equilibrium analysis suggests that one might expect to observe empirically less work (and less success) on projects when they are first initiated then when they are continued.[7]

---

[7] Rivkin et. al. (2006) describes how a CEO-driven initiative to move Whirlpool to become an innovation-oriented company had slow initial progress amid widespread employee skepticism and only picked up momentum after the firm showed a deep commitment to building an innovation culture. Such a pattern of work

Finally, we explore the implications of persistence in a complete information environment when effort by either of two decentralized managers (research teams) may lead to a successful project with subsequent outcomes and payoffs that apply equally to both managers. We find that the presence of a second manager introduces free-riding incentives which results in overall less success on projects and less persistence by leaders when the cost of effort is high relative to the expected value of projects.[8]

The complete and incomplete information analyses underscore the importance of the underlying distribution for the leader and manager choices. The intuition from the R&D project setting may carry over to other settings involving continued investment in a particular person (junior faculty, consulting or legal firm associations, professional athletes) or in existing market products. Further, with some changes in interpretation, our persistence model arguably offers insights into attempts by leaders to inspire and motivate subordinate efforts by communicating a vision for the organization's future.

A primary focus of the economic literature regarding innovation management is on the impact of a principal's incentive schemes on an agent's innovation activities. Manso (2011), for example, builds a two-period principal-agent model in which agents have a choice in each period of pursuing innovation through either an exploitation or an exploration project. The exploitation project has a known probability of success. If the agent pursues the ex-

---

is consistent with a "wait and see" attitude as described by Bower and Gilbert (2005) and is common in other situations such as government bureaucrats' reactions to initiatives proposed by a new and inexperienced political appointees (Heclo 1977).

[8] See Khanna et. al. (2018) for a discussion about how high interdependence across research programs decreases terminations of invention projects because decision makers are boundedly rational. In contrast, we focus on how free-riding incentives among interdependent project teams reduces effort, with equilibrium consequences for the persistence threshold decision and project success outcomes.

ploration project and fails, the agent will learn more about the probability of success which then impacts the subsequent choice. Incentivizing exploration (as opposed to exploitation) is encouraged by the principal through a commitment to an incentive scheme for the agent that allows for failure and takes into account the performance path. Hellmann and Theile (2011) analyze a setting in which an agent divides effort between a standard (exploitation) task and an innovative (exploration) task. The principal offers incentives through an ex ante contract based on performance on the standard task with an innovative task success resulting in ex post contracting. The agent privately observes the potential value of an innovative outcome before deciding on effort. The incentives on the standard task change the agent's willingness to engage more on the innovation task.[9]

We differ from these papers by focusing on the impact of termination of projects on which employees are assigned rather than the impact of the structure and level of direct incentive schemes. Like Manso and Hellmann and Theile, search is a critical part of the story, but in our model innovative search is constrained by the leader's choice of projects: search in our model is not the agent's choice, it is the leader's choice. We see our approach as emphasizing leadership rather than incentives. By focusing on leader-driven search we also open up the question of a leader's private information about the value of a project to the firm and the impact of private information on effort. At a high level, this private information about value can be thought of as trying to align subordinates with the leader's vision of the value of particular projects or even (though outside the scope of this disser-

---

[9] See, also, Bergemann and Hege (2005) who explore a principal-agent financial relationship between a venture capitalist and an entrepreneur over time. Their model involves a single project and focuses on the entrepreneur's incentives to divert funds away from the primary project whose prospects and timetable are uncertain.

tation chapter) the portfolio of innovation projects. Persistence as a policy instrument can do this, arguably better than incentive schemes, though the two are not, of course, mutually exclusive instruments.

Treating the principal rather than the agent as having private information is also central to Hermalin's (1998) paper on leadership. There, Hermalin examines a model with a principal who is informed about the value of a project, a project's ex post value depends on the team efforts of both the principal and the agent, and the principal wishes to induce effort on the part of an agent. Hermalin shows how "leading by example"—having the principal exert observable effort—can effectively signal the principal's information and induce agent effort. Persistence in our model plays a similar signaling role, but the mechanism relies on the equilibrium implications of optimal search (i.e., continue to invest in the current project or move to a new project?).[10]

An important component of our analysis involves search for projects with superior value. The search process we employ is a simplified version of a "bandit" problem in which a decision maker optimally decides whether to continue to pull one lever (invest again in the current project) or switch to a new lever (switch investment to a new project).[11] While we share an interest in understanding the organizational decision to switch, our emphasis is different than that of the search literature because we seek to understand the interaction between search and agent effort incentives. Hence, we chose not to incorporate processes

---

[10] Aghion and Tirole (1997) investigate how allocating project selection authority to either the principal or the agent affects the incentives of a principal and an agent to invest in learning the value of several candidate projects. Given our assumption that the leader (principal) learns the value of a project once drawn from the distribution of projects, our model cannot be used to address this problem. See also Rantakari (2012).

[11] See Bergemann and Välimäki (2002) for a review of these problems.

such as learning or the value of real options into our analysis.[12] Such processes would also generate persistence. For example, Halac and Liu (2016) analyze long-term contracts in the presence of agents of different abilities with learning about the quality of a project. The duration of experimentation—the optimal stopping rule for experimentation–is a form of persistence motivated by learning incentives. Among other differences, their model focuses on experimentation with a given project while we emphasize search over a set of projects. Search over a set of projects means that the principal controls what projects the agent may work on. In particular, the principal can replace the existing project with a new project.[13]

In the next section, we describe the formal model and establish some benchmarks for the equilibrium analysis. Sections 3 and 4 describe our complete information results regarding equilibrium persistence under different assumptions about the underlying distribution of future projects. Section 5 analyzes persistence under incomplete information where the leader has private information about the value of the project. Section 6 examines the value of persistence in coordinating the actions of two subordinates (or research teams) and Section 7 discusses some general implications of our analysis, notes some limitations, and concludes.

---

[12] This is true except in a very limited sense in the incomplete information model where the principal's knowledge of the value of a project is signaled to the manager. There is no mechanism for learning about future project draws in the model. There is a substantial literature focusing on learning and real options as applied to R&D project termination (e.g., Brunner and Zinner (2008) for a review).

[13] Halac and Liu (2016) argue that a solution to problems without adverse selection is to sell the value of the project to the agent which achieves the first best. Presumably, this is because there are not contracting frictions associated with agent access to needed assets, teams, and intellectual property. We assume this type of arrangement is not possible in a corporate setting and build our model to implicitly reflect organizational constraints that limit the range of a leader's contract arrangements. See Section 7.

4.2 Model

Our primary concern is to explore the persistence instrument and its impact on manager effort. This problem is most interesting when there is an incentive conflict between the person making the persistence decision and the person exerting effort. While the incentive conflict problem is standard, a key feature of our model is that it separates project choice from effort and the project is chosen by the leader.[14]

We model the persistence decision as a form of project selection which is controlled by the leader (principal). The leader can choose to persist with the current (as yet unsuccessful) project or can search by initiating a new project drawn from a distribution of projects which differ only by the value of the project given success. While the value of a specific project is unknown when initially selected, the characteristics of the underlying distribution of projects in known. This structure is consistent with the description of decision making offered by many management scholars (e.g., Bower and Gilbert 2005) in which the importance of higher-level managerial control over project selection is noted.[15]

The choice and search structure of our model is integrated into a model of multi-period innovation. Guler 2007 and the CBO 2021, among others, describe risky R&D project choices as being characterized by the importance of lack of payoffs until late in the project

---

[14] This assumption differs from the route taken with the project experimentation literature in economics which is concerned with how incentive structures affect an agent's choice of whether to explore a more uncertain project or exploit a less uncertain project. In our model the leader (principal) selects the R&D project and the subordinate (agent) chooses effort. Effort in our model could be interpreted as choosing between explore or exploit options, but the key assumption is that the explore project(s) are chosen by the principal.

[15] Of particular relevance is the resource allocation process (RAP) literature pioneered by Bower (1970). See Section 7.

cycle, the opportunity cost of the resources used, and uncertainty about net expected value of the project which is reduced over time. Our model of the innovation process incorporates the first two features, but largely suppresses the last.

The innovation part of the model captures the inherent uncertainty of the process which is necessary for persistence to matter. Ex ante, a new project has unknown value, but becomes known to the leader once the project is drawn. Under complete information the manager also learns the project value. Hence, learning about the project value is exogenous and is not the result of choice. In each period, manager effort results in a probability of success that is independent of prior effort. The innovation and information about project value assumptions are quite strong, but have the advantage of isolating the impact of persistence on effort. That is, the persistence decision is not about gaining further knowledge for the leader regarding the expected value of a successful project. It is about manager effort and the relative expected values of the current project and a new project draw from the distribution of possible projects.[16]

Formally, our base complete-information model involves a leader L and a manager M who make decisions on a sequence of projects. Time is discrete, indexed by $t = 1, 2, \ldots$, with an infinite horizon. The value of each project is an independent draw from a common distribution F with density f, support $[0, \infty)$ and mean μ. Both of L and M observe u_t, the value of the available project for period t (later, under incomplete information only L observes u_t). First, M decides whether to expend effort on the project, $e_t \in \{0, 1\}$. This

---

[16] Some studies of sequential investment decisions in risky R&D projects focus on information asymmetries noting that the subordinate (agent) will have better knowledge about the technical aspects of the project. We focus on the impact of the leader's superior knowledge of the commercial attractiveness of the project and how the project fits in with the overall portfolio of projects and activities of the firm.

choice is private and L does not observe the effort level. If M works, $e_t = 1$, the project succeeds with probability $p \in (0,1)$. The project fails if $e_t = 0$. The expected period-t payoffs are $pu_t - c$ for M and pu_t for L if M works, and 0 for M and L if M does not work on the project (henceforth, "shirk"). The reward is common, but only M incurs the effort cost $c > 0$.[17]

If the project for period t is a success, then the project for period $t+1$ is a new draw from the distribution F. In the event of a failure, however, L decides whether to persist with the project ($\rho_t = 1$) and try again in period $t+1$, so that $u_{t+1} = u_t$, or to abandon the project ($\rho_t = 0$) for a new draw from F.

Each player maximizes the expected present discounted value of their period payoffs, and $\delta \in (0,1)$ is the discount factor. Reflecting the stationarity in the structure of the model, we focus on stationary equilibria where, by definition, the effort choice of M and the persistence choice of L, denoted by $e^*(u)$ and $\rho^*(u)$, respectively, depend only on the value u of the current project.

Definition 2. *A Persistence Equilibrium (PE) is a stationary equilibrium in which there is a* $\bar{u} \geq 0$ *such that (i) in period* t *in the event of a failure, the leader* L *abandons the project if* $u_t < \bar{u}$ *and persists if* $u_t \geq \bar{u}$*; (ii) in period* t*, the effort choice* $e^*(u_t)$ *of the manager* M *is optimal; and (iii) in period* t*, the persistence choice* $\rho^*(u_t)$ *of the leader* L *is optimal.*

We will see shortly that a PE is characterized by two thresholds, one for effort by M and one for persistence by L. While the threshold for the persistence choice by L is part

---

[17] The simplest interpretation is that a project has a total value of $2u_t$ with equal shares for L and M. A natural extension is to allow for unequal division of the total reward.

of the definition, a threshold for the effort choice by M emerges as a consequence of the incentives created by the persistence threshold.

Benchmarks A few benchmarks will be helpful for understanding the equilibrium results. Consider the simplest possible setting where there is only one period so that there is no persistence decision and M makes a one-time static effort choice. If the project draw is u, then M will find it optimal to work when $pu \geq c$ and to shirk (not work on the project) otherwise. This yields static payoffs of $pu - c$ for M and pu for L. Of course, since L does not incur the cost of effort, we have a wedge between the players and L would prefer that M works on all positive value projects. The ex ante payoff for M is then

$$\int_{c/p}^{\infty} (pu - c) f(u) du = p \int_{c/p}^{\infty} u f(u) du - c[1 - F(c/p)]$$

and the ex ante payoff for L is, of course, the value expectation ignoring the effort cost.

To assess equilibrium dynamics, it will be helpful to refer to an infinite repetition of the above static one-period interaction where there is a new draw from F every period (i.e., no persistence decision by L). Effort by M will follow the c/p cutoff threshold and payoffs are simply the static payoffs divided by $1 - \delta$ to measure discounted value.

4.3 Equilibrium Persistence with a Bright Future

We begin the analysis of equilibrium persistence by considering the possibility of a PE in which the $\bar{u}$ threshold is strictly greater than c/p, the benchmark effort cutoff in the static setting. Define this as a "high-$\bar{u}$" equilibrium. In such an equilibrium, L will choose to abandon projects that have a strictly positive, potentially significant, value in a static setting. We begin with effort incentives for M, develop the value functions, and then move to incentives for L. With these results in hand, we present the equilibrium characterization result.

Effort incentives in a high-$\bar{u}$ equilibrium are simple. Whenever the current project has value $u > c/p$, the current period expected payoff of $pu - c$ justifies the effort. Since providing no effort in the future is always an option, the value of the future cannot be negative. Hence, we have the basic result in

Lemma 11. *In any PE with $\bar{u} > c/p$, we have effort $e^*(u_t) = 1$ if and only if $u_t \geq c/p$. (See Appendix for all proofs.)*

Thus, effort in a high-$\bar{u}$ equilibrium follows the same pattern as in the static outcome. Importantly, however, even though M works when the project value u_t is between c/p and $\bar{u}$, this does not lead L to persist. That is, if such a project is unsuccessful, it will be terminated by L in favor of a new draw from F. Intuitively, this outcome is driven by the expectation of a "bright future" meaning that there is sufficient expected value from a new draw to justify termination.

The structure of value functions for M and L follow directly from the effort pattern in Lemma 11. For M and L respectively, we let $v(u)$ and $v^L(u)$ denote the interim value functions when the current project has payoff u, while $V(\bar{u})$ and $V^L(\bar{u})$ denote the ex ante (with respect to the next project draw from F) value functions for an equilibrium with threshold $\bar{u}$. In all cases, we can move from the value function of M to that for L simply by eliminating the cost term. Begin with the interim value and suppose the current project has $u < \bar{u}$. Since this project will be terminated, we see that for $u < \bar{u}$

$$v(u) = (pu - c)e^*(u) + \delta V(\bar{u})$$
$$v^L(u) = pue^*(u) + \delta V^L(\bar{u})$$

relates the interim and ex ante values. Recall that effort is given by Lemma 11 for u above and below c/p. In contrast, when $u \geq \bar{u}$ and L will persist with an unsuccessful current project, we have

$$v(u) = (pu - c)e^*(u) + p\delta V(\bar{u}) + (1-p)\delta v(u) \qquad \Rightarrow$$
$$v(u) = \frac{pu - c + p\delta V(\bar{u})}{1 - \delta(1-p)}$$

and similarly

$$v^L(u) = \frac{pu + p\delta V^L(\bar{u})}{1 - \delta(1-p)},$$

where we have made use of Lemma 11 to set $e^*(u) = 1$ as, in this case, $u > \bar{u} > c/p$. Value functions are in each case a weighted sum of current and future payoffs. Note that the weights are shifted in termination states versus persistence states. Current payoff flows are weighted less heavily in termination states, since we always have $1 < 1/[1-\delta(1-p)]$, while future payoffs are more important in persistence states, since $\delta > p\delta/[1-\delta(1-p)]$.

With the interim value functions from above, we can determine the ex ante value functions. Integrating across u draws, we have for the manager M

$$\begin{aligned} V(\bar{u}) &= \int_0^\infty v(u)f(u)du \\ &= \int_0^{\bar{u}} [(pu-c)e^*(u)+\delta V(\bar{u})]f(u)du + \int_{\bar{u}}^\infty \frac{pu-c+p\delta V(\bar{u})}{1-\delta(1-p)} f(u)du \end{aligned}$$

as follows from the interim value for u above and below the threshold. Simplifying and applying Lemma 11 for the effort level, we can solve for the ex ante value to find

$$V(\bar{u}) = \frac{\int_{c/p}^\infty (pu-c)f(u)du - \delta(1-p)\int_{c/p}^{\bar{u}}(pu-c)f(u)du}{(1-\delta)[1-\delta(1-p)F(\bar{u})]}$$

for the manager M while repeating the same exercise yields

$$V^L(\bar{u}) = \frac{\int_{c/p}^\infty puf(u)du - \delta(1-p)\int_{c/p}^{\bar{u}} puf(u)du}{(1-\delta)[1-\delta(1-p)F(\bar{u})]}$$

for the leader L. As with the interim case, ex ante values only differ between M and L by the effort cost incurred by M, but cost also matters to L since it determines the full range of projects over which M will expend effort.

Turning to incentives, consider the optimal choices by M and L in each state. The situation for M is simple. Incentive compatibility (IC) in a persistence state where $u \geq \bar{u}$ requires that work yields a higher payoff than not working:

$$pu-c+p\delta V(\bar{u})+(1-p)\delta v(u) \geq 0+\delta v(u)$$

since not working implies a failure and then persistence by L. Simplifying with $v(u)$ from above, this reduces to

$$pu - c + p\delta V(\bar{u}) \geq 0 \tag{48}$$

which holds for all $u \geq \bar{u}$ if it holds at the threshold $\bar{u}$. Since a high $\bar{u}$ is above the static work threshold of c/p and the ex ante value is positive, M will choose to work in persistence states. IC for M in termination states holds trivially. As noted above with Lemma 11, success or failure and, hence, work effort has no impact on the future payoff since the current project will be terminated. IC in this case then reduces to working when u is above c/p and not working when below.

Now, consider the substantive question of the incentive for persistence by L. When $u < c/p$ we know that M does not work and the current project must fail. Thus, persistence would simply reproduce the return of 0 in the next period. Formally, incentive compatibility for L in this state requires that termination dominates persistence, or $\delta V^L(\bar{u}) \geq \delta^2 V^L(\bar{u})$, which clearly holds. Termination must also be the equilibrium choice over the middle range of project states where $c/p < u < \bar{u}$. Since M works in these states, persistence by L will yield $p[u + \delta V^L(\bar{u})] + (1-p)\delta v^L(u)$ while termination will yield $pu + \delta V^L(\bar{u})$. Simplifying with $v^L(u)$ and comparing, we see that termination dominates persistence if

$$\frac{pu}{1-\delta} \leq V^L(\bar{u}).$$

Intuitively, persisting in repeated attempts with the project u, the discounted value on the left, is not sufficient in value relative to terminating for a draw on a new project. Over the high range of states, persistence must be the equilibrium choice and incentive compatibility, by the reverse of the middle range comparison, requires that $pu/(1-\delta) \geq V^L(\bar{u})$ when $u > \bar{u}$. Thus, incentive compatibility for L across all u values requires that the threshold satisfies

$$\frac{p\bar{u}}{1-\delta} = V^L(\bar{u}). \tag{49}$$

Combining the earlier expression for the ex-ante value function for L with this incentive compatibility condition, we are able to characterize the threshold.

The following proposition establishes existence and uniqueness for a high-$\bar{u}$ persistence equilibrium.

Proposition 20. *Suppose that*

$$\frac{c}{p}\left[1-\delta(1-p)F\left(\frac{c}{p}\right)\right] < \int_{c/p}^{\infty} uf(u)du \tag{50}$$

holds. Then there exists a unique PE in which the threshold $\bar{u}$ satisfies $\bar{u} > c/p$. If (50) does not hold, then there does not exist a PE with $\bar{u} > c/p$.

The equilibrium threshold $\bar{u}$ is determined by

$$\bar{u}[1-\delta(1-p)F(\bar{u})] + \delta(1-p)\int_{c/p}^{\bar{u}} uf(u)du = \int_{c/p}^{\infty} uf(u)du. \tag{51}$$

Note that, on the right, we have the expected value of a project draw with the lower limit reflecting the equilibrium effort by M. This provides a reference value for L with respect to terminating a current project. On the left, we have a reference value for L with respect to persisting with a project where the draw is exactly $\bar{u}$. The threshold is determined by balancing these options. For a higher u, persisting is optimal for L. For a lower u, terminating is optimal.

How does the threshold relate to future prospects for projects? Given the different integral limits in (51), it is helpful to sort out how a simple feature such as the mean for project draws under F will influence the equilibrium.

Corollary 3. *Suppose that*

$$\frac{c}{p}[2-\delta(1-p)] < \mu$$

where $\mu = \int_0^\infty uf(u)du$ is the mean project draw. Then (50) holds and there exists a unique PE where $\bar{u} > c/p$.

The intuition is clear. With a “bright future,” meaning that the mean for project draws μ is sufficiently large, we necessarily have a PE with a high threshold. It is also worth noting that the sufficient condition in Corollary 3 is more easily satisfied with a larger p. With a higher probability for success, M works over a larger range of projects, as c/p falls.

When the future is bright, L, in equilibrium, will terminate projects that may be well above c/p in value. For example, terminating a sizeable u with a high p may require L to provide “justification” to investors and outside observes (e.g., industry and financial press). In practice, that could be done through a “soft” termination by starving a project of resources rather than officially killing it. Later successes will emerge to cover such misfires. Note as well that IC for L implies that the leader will never choose to restart a project that was previously terminated.

4.4 Equilibrium Persistence with Dim Prospects

We now turn to the possibility of equilibrium persistence with a threshold that is below the static threshold for effort, $\bar{u} \le c/p$. Define this as a "low $\bar{u}$" equilibrium. Now, in contrast to the high-$\bar{u}$ equilibrium, L will never choose to abandon any project that has a strictly positive value in a static setting. More intriguingly, L will also choose to persist with projects that have a negative value from a static viewpoint, those where $\bar{u} \le u \le c/p$. Moreover, it must be that M chooses to provide effort on these projects. As before, we progress from effort incentives for M, to the value functions, and then to incentives for L. The equilibrium characterization result is then developed.

Effort incentives in a low $\bar{u}$ equilibrium are more subtle than in the high $\bar{u}$ case. This is because dynamic incentives are necessary for creating an incentive to work on projects when u is below c/p. A direct implication of the PE definition for work effort in this case is

Lemma 12. *In any PE with $\bar{u} \le c/p$, we have effort $e^*(u_t) = 1$ if and only if $u_t \ge \bar{u}$.*

Effort in a low-$\bar{u}$ equilibrium must thus extend below the static benchmark. Now, in contrast to the high-$\bar{u}$ case, L will persist with a range of projects that are below the static benchmarks. Intuitively, in contrast to a "bright future," this low-range persistence must be justified by an expectation that a new draw does not justify termination. Thus, it is only through effort and a success that the parties will move on from the current project.

Turning to the value functions, for a project that will be terminated, $u < \bar{u}$, we have

$$v(u) = (pu - c)e^*(u) + \delta V(\bar{u}) = \delta V(\bar{u})$$

$$v^L(u) = pue^*(u) + \delta V^L(\bar{u}) = \delta V^L(\bar{u})$$

to relate the interim and ex ante values. Recall that effort is 0 in this case, as in Lemma 12. When $u \geq \bar{u}$ and L will persist with the current project in the event of failure, we have, as with the high-$\bar{u}$ case,

$$\begin{aligned} v(u) &= (pu - c)e^*(u) + p\delta V(\bar{u}) + (1-p)\delta v(u) \qquad \Rightarrow \\ v(u) &= \frac{pu - c + p\delta V(\bar{u})}{1 - \delta(1-p)} \end{aligned}$$

and similarly

$$v^L(u) = \frac{pu + p\delta V^L(\bar{u})}{1 - \delta(1-p)},$$

where we have made use of Lemma 12 to set $e^*(u) = 1$ as, in this case, $u > \bar{u}$.

Turning to the ex ante values, we integrate the interim values over u draws and simplify to find

$$V(\bar{u}) = \frac{\int_{\bar{u}}^{\infty}(pu - c)f(u)du}{(1-\delta)[1 - \delta(1-p)F(\bar{u})]}$$

for the manager M while repeating the same exercise yields

$$V^L(\bar{u}) = \frac{\int_{\bar{u}}^{\infty} puf(u)du}{(1-\delta)[1 - \delta(1-p)F(\bar{u})]}$$

for the leader L.

Now consider incentives for optimal choices by M and L in each state. Incentive compatibility for M when $u < \bar{u}$ is trivial since M does not work and L terminates the project. Termination fixes the future payoff at $V(\bar{u})$, and working would only yield a negative expected current payoff of $pu - c < 0$ since $u < \bar{u} < c/p$. For $u \geq \bar{u}$, work must be preferred to shirking on the project and this requires $pu - c + p\delta V(\bar{u}) \geq 0$, as shirking implies project failure and then persistence by L; note that this logic only requires knowing that $u \geq \bar{u}$, so we arrive at (48) just as with the high-$\bar{u}$ threshold. Combining with $V(\bar{u})$ from above for the low-$\bar{u}$ threshold, noting that if it is optimal to work in state $\bar{u}$ then it is optimal to work when $u > \bar{u}$, incentive compatibility for M requires that

$$p\bar{u} - c + p\delta\left[\frac{\int_{\bar{u}}^{\infty}(pu - c)f(u)du}{(1-\delta)[1-\delta(1-p)F(\bar{u})]}\right] \geq 0 \tag{52}$$

holds at $\bar{u}$. We then have

Lemma 13. *IC for M holds for an interval of threshold values* $[\bar{u}_M, c/p]$ *where* $0 \leq \bar{u}_M < c/p$. *If* $\frac{c}{p}[1-\delta(1-p)] < \delta p\mu$, *then* $\bar{u}_M = 0$.

Thus, there is always a range for the threshold below c/p such that IC for M will hold for any $\bar{u}$ in this range. Intuitively, this is because the work incentive for M in (48) is the sum of the current expected return, $pu - c$, and the term reflecting the value of future projects, $p\delta V(\bar{u})$. Because the value of future projects is strictly positive, there is leverage to offset a negative current expected return for a given $\bar{u}$. Hence, a threshold below c/p can always support effort incentives and M is willing to provide effort on projects where u is below c/p but above the threshold $\bar{u}$.

For another perspective, consider shirking when $\bar{u} < u < c/p$. In the stationary context of a PE, shirking today implies a failure today and, consequently, a zero payoff today from the current project. Moreover, L will persist with current project. With stationarity, this means that shirking simply pushes the interim payoff of $v(u)$ one period into the future. Hence, discounting creates a necessary wedge that implies work must dominate if we have a current project draw u that is only slightly below c/p. Essentially, persistence by L induces M to work on lower value projects because that is the only way for M to move on to a new project draw. Finally, the sufficient condition in Lemma 13 reveals that IC for M may impose no restriction at all on the threshold. This occurs when the static threshold, c/p, is a fraction of the project mean, μ, where the fraction $\delta p/[1-\delta(1-p)]$ is always below 1.

Now consider IC for L. In contrast to the equality condition for IC in (49) in the high-$\bar{u}$ case, the persistence incentive is an inequality in the low-$\bar{u}$ case. First, note that IC for L reduces trivially to $\delta V^L(\bar{u}) < V^L(\bar{u})$ when $u < \bar{u}$ since persistence simply replicates the no-effort state and failure outcome. Thus, when the threshold $\bar{u}$ is below c/p, L will always terminate when $u < \bar{u}$. When $u \geq \bar{u}$, the same logic as in the high-$\bar{u}$ case leads to

$$\frac{p\bar{u}}{1-\delta} \geq V^L(\bar{u}) \qquad \Leftrightarrow$$

$$\bar{u} \geq \frac{\int_{\bar{u}}^{\infty} uf(u)du}{1-\delta(1-p)F(\bar{u})}. \tag{53}$$

Persistence incentives for L contrast with the effort incentives for M. First, consider whether L will ever be willing to persist when a project has a very low value. We know from Lemma (13) that M is willing to work all the way down to $u = 0$ when the sufficient condition holds. But this is never the case with L and persistence. Consider IC for L in (53): a threshold of $\bar{u} = 0$ would require a negative mean for projects. In turn, this means L will never be willing to persist for projects near 0 in value. Second, consider a threshold that is below but close to c/p. As we know from Lemma 13, M will necessarily have an incentive to provide effort in this case. As we see next, this does not carry over to the persistence incentive for L. In other words, the existence of a low-$\bar{u}$ PE hinges on whether L will choose to persist with low-value projects, not on whether M will provide effort.

In analogy with the "bright future" interpretation of the condition on the mean μ for a high-$\bar{u}$ PE, the existence of low-$\bar{u}$ PE requires that the same inequality goes in the other direction, leading to the interpretation that the future has "dim prospects."

Proposition 21. *Suppose that*

$$\frac{c}{p}\left[1-\delta(1-p)F(\frac{c}{p})\right] > \int_{c/p}^{\infty} uf(u)du \tag{54}$$

holds. Then there exists a $\bar{u}_0 \in (0, c/p)$ *such that for any* $\bar{u} \in (\bar{u}_0, c/p)$ *there exists a* PE *with threshold* $\bar{u}$.

The condition (54) directly implies that IC holds for L in an interval below c/p. Note that (54) is simply the reverse of (50), the existence condition for a PE with high-$\bar{u}$. Thus, the two types of equilibria are mutually exclusive, and the condition on the mean determines which type will exist. While IC for M can potentially extend all the way down to $\bar{u} = 0$, meaning that M can be induced to work in every state, the same is not true of IC for L. This is why $\bar{u}_0 > 0$ holds.

Proposition 21 implies that, when (54) holds, we will have multiple equilibria in that each $\bar{u}$ threshold above $\bar{u}_0$ and below c/p will constitute a PE. In contrast, a high-$\bar{u}$ equilibrium, which exists when (54) fails, is always unique. The reason for this difference lies with effort incentives and how they interact with persistence incentives. In a high-$\bar{u}$ equilibrium, the threshold is above c/p but M provides effort whenever the project value u is above c/p. Thus, the threshold is neither determining when M works nor are the work incentives for M pinning down the threshold. All that matters for effort here is that the threshold is above the static benchmark of c/p. In a high-$\bar{u}$ equilibrium, the threshold is determined by IC for L according to (49), and L is indifferent between persistence and termination when $u = \bar{u}$. Persistence is optimal when u for the project is higher and termination is optimal when u is lower. There is only one threshold that is consistent with persistence for projects above and termination for projects below this threshold.

The range of equilibrium thresholds with low-$\bar{u}$ equilibria is suggestive of an implicit coordination between M and L in that effort and persistence must reinforce each other in the range below c/p. The first point to note is the contrast in effort. With a high-$\bar{u}$ equilibrium, M provides effort on projects where u is below $\bar{u}$; when u is above c/p the positive current return justifies effort even though L will terminate the project. This never occurs in a low-$\bar{u}$ equilibrium. For u below c/p, M will never provide effort if u is below the $\bar{u}$ threshold: effort on such a project has a negative current return and L will terminate. In other words, there is no effort when termination is expected. This difference also carries over to L. With a high-$\bar{u}$ equilibrium, L would terminate projects even though M provided effort, as when u was below $\bar{u}$ but above c/p. This also never occurs in a low-$\bar{u}$ equilibrium. Without effort on project, L will always terminate.

Another way to see the contrast is to note that M and L each have one-sided IC constraints, in (52) and (53), in a low-$\bar{u}$ equilibrium. For any threshold above the lower bound of $\bar{u}_0$, both IC constraints are slack (strict inequality at the threshold). The lower bound $\bar{u}_0$ is determined by which of (52) and (53) reaches equality first. For example, when the sufficient condition in Lemma 13 holds and, hence, (52) is always strict, $\bar{u}_0$ is determined by IC for L.

In analogy, with the high-$\bar{u}$ equilibrium, it is helpful to identify a relaxed version of the existence condition for a low-$\bar{u}$ equilibrium. In contrast to a "bright future" (high mean) that supports a high threshold, we find that "dim prospects" support a low threshold for persistence:

Corollary 4. *A sufficient condition for the existence of a low-$\bar{u}$ equilibrium is*

$$\frac{c}{p} > \frac{\mu}{1-\delta(1-p)},$$

which implies that (54) holds.

As before, it is the relative and not the absolute size of μ that matters. Thus, if we have a high static threshold for effort, either because c is large or p is low, then μ will be in the “dim prospects” category.

The high-$\bar{u}$ and low-$\bar{u}$ PE analyses underscore the importance of the distribution of project values F (e.g., bright or dim prospects) relative to the expected value of effort $pu-c$ for which projects are continued. With bright prospects, more projects are terminated on average and the manager will work only if the static expected value of working is positive. With dim prospects, fewer projects are terminated and the manager works on continued projects even though the static expected value of working is negative. In terms of empirical predictions, these results suggest that firms facing a poor value distribution of projects will successfully complete more projects than firms with a good distribution because, on average, those firms will persist more frequently and managers will work more often. Firms with a good distribution will terminate projects more frequently and have fewer successfully completed projects, but the projects completed will be of greater value.[18]

4.5 Persistence and Incomplete Information

We now consider settings in which the leader has better information about the value of a project. Thus, when deciding whether to persist with or to terminate a project, the leader bases that decision on private information. In contrast, the manager makes the effort decision with more limited information. Importantly, however, the choice of the leader to persist may influence the beliefs of the manager regarding project value and, therefore, the choice to provide effort.

---

[18] Another factor impacting the use of persistence is whether the firm is experiencing financial distress which might manifest as a higher weight on current period payoffs, i.e., a smaller δ in our model.

Formally, the only change we make to the model is that we now assume only L observes the project value u_t. While M does not observe u_t, M does observe the persistence decision of L. Thus, when choosing effort, M knows whether it is the initial stage of a project or whether L has persisted and, hence, if the current project is the same one as in the previous period. We let $e_i \in \{0,1\}$ and $e_p \in \{0,1\}$ denote the effort choice in an initial state and in a persistence state, respectively. We define an equilibrium by

Definition 3. *A Persistence Equilibrium (PE) under incomplete information is a stationary (Perfect Bayesian) equilibrium in which there is a* $\bar{u} \geq 0$ *such that (i) in each period in the event of a failure, the leader* L *abandons the project,* $\rho^*(u_t) = 0$*, if* $u_t < \bar{u}$ *and persists,* $\rho^*(u_t) = 1$*, if* $u_t \geq \bar{u}$*; (ii) if period* t *is the initial stage with a project, then the effort choice* e_i^* *of the manager* M *is optimal; (iii) if period* t *is a persistence stage with a project, then the effort choice* e_p^* *of the manager* M *is optimal; (iv) the persistence choice* $\rho^*(u_t)$ *of the leader* L *is optimal.*

Thus, a PE will consist of three elements. First, we have the cut-off level of $\bar{u}$ for the persistence decision by L. The other two elements are the effort levels for M in an initial project state, e_i^*, and in a persistence state, e_p^*. In an initial state, the beliefs of M regarding project value are that u is drawn from F. Thus, with respect to the current period flow payoff from effort in an initial project state, M expects $p\mu - c$ as the return to effort, where μ is the (unconditional) mean of the project distribution F. In a persistence state, M updates beliefs to account for the persistence choice of L. Since L follows a cut-off rule, in a persistence state M knows that values below $\bar{u}$ are no longer relevant. The updated belief of M is conditioned on the event $\{u \geq \bar{u}\}$ so that project value is distributed as $f(u)/[1-F(\bar{u})]$. Then, instead of basing the effort choice on the mean, μ, under the updated belief M will expect a current period flow payoff from effort of $p\bar{\mu}-c$, where $\bar{\mu}=\int_{\bar{u}}^{\infty} uf(u)du/[1-F(\bar{u})]$ is the conditional mean. Of course, $\bar{\mu}$ is strictly increasing in $\bar{u}$ and satisfies $\bar{\mu}>\bar{u}$. Thus, a PE can be summarized by a triple $(\bar{u}, e_i^*, e_p^*)$.

There are three possible configurations for a PE. One possibility is a full-effort PE in which M provides effort in both states, with $e_i^*=e_p^*=1$. Second, there may be a PE with effort in one state and shirking in the other. As we shall see, this pattern only occurs with $e_i^*=0$ and $e_p^*=1$, so that persistence overcomes initial shirking in equilibrium by shifting beliefs about the project. Finally, we may have no effort in equilibrium, $e_i^*=e_p^*=0$, a type of organizational failure under incomplete information.

The value functions in an equilibrium with efforts e_i^* and e_p^* by M and a persistence threshold of $\bar{u}$ by L are straightforward to construct. For L, we will need the interim value $v^L(u)$ and the ex ante value V^L. For the interim case, in equilibrium, when $u<\bar{u}$ it must be that L terminates the project in the event of failure in the current period. Hence, we have $v^L(u)=e_i^*pu+\delta V^L$ when $u<\bar{u}$. Note that this interim value applies to initial states (and to persistence states if we are off equilibrium, for instance, if L has deviated). For persistence states where $u \geq \bar{u}$, in equilibrium, L will persist after failure and we have

$$v^L(u)=e_p^*pu+e_p^*[\delta pV^L+\delta(1-p)v^L(u)]+(1-e_p^*)\delta v^L(u) \qquad \Longrightarrow$$

$$v^L(u)=\frac{e_p^*[pu+\delta pV^L]}{1-\delta(1-pe_p^*)} \qquad \text{for } u \geq \bar{u}.$$

The ex ante payoff for L is then calculated as

$$V^L = \int_0^\infty v^L(u)f(u)du = \int_0^{\bar{u}} [e_i^* pu + \delta V^L]f(u)du + \int_{\bar{u}}^\infty [e_i^* pu + \delta v^L(u)]f(u)du$$
$$= e_i^* p\mu + \delta V^L F(\bar{u}) + e_p^* \delta \frac{p \int_{\bar{u}}^\infty uf(u)du + \delta p V^L[1 - F(\bar{u})]}{1 - \delta(1 - pe_p^*)}$$

which simplifies to

$$V^L = \frac{e_i^* p\mu[1 - \delta(1 - pe_p^*)] + e_p^* \delta p \bar{\mu}[1 - F(\bar{u})]}{(1 - \delta)[1 + e_p^* \delta p - \delta F(\bar{u})]}.$$

Turning to the value functions for M, the new element is on the information side. M does not observe the project value u. Instead, M only knows whether it is the initial state with a project or whether it is a persistence state. Because of this, there is no distinction between the ex ante value for M and the interim value in the initial state. So, let us begin with v^M, the interim value for the persistence state. Knowing that L has persisted, a choice not to provide effort will yield a flow payoff of zero and, since failure is the certain outcome, a subsequent choice to persist by L. Choosing to provide effort will yield a flow payoff from the current chance for success and, depending on that outcome, either a new draw or the persistence response by L. Thus

$$v^M = e_p^*[(p\bar{\mu} - c) + \delta p V^M + \delta(1 - p)v^M] + (1 - e_p^*)[0 + \delta v^M] \qquad \Longrightarrow$$

$$v^M = e_p^* \frac{p\bar{\mu} - c + \delta p V^M}{1 - \delta(1 - pe_p^*)}.$$

Note that the interim value in the persist state is based on the conditional mean $\bar{\mu}$, reflecting the updated belief of M regarding project value while, in contrast, the interim value for L is based on the actual u project draw. The other difference, of course, is that M incurs the effort cost. Next, for the ex ante value V^M, we must calculate the path following the current effort choice and account for the likelihood of whether the initial stage involves a project with value above or below the persistence cutoff. Thus,

$$
\begin{aligned}
V^M = {} & e_i^*[(p\mu - c) + \delta p V^M + \delta(1-p)[F(\bar{u})V^M + (1-F(\bar{u}))v^M]] \\
& + (1-e_i^*)[0 + \delta F(\bar{u})V^M + \delta(1-F(\bar{u}))v^M]
\end{aligned}
$$

which, upon substituting for v^M, simplifies to

$$
V^M = \frac{e_i^*(p\mu - c)[1-\delta(1-pe_p^*)] + e_p^*(p\bar{\mu} - c)\delta(1-pe_i^*)[1-F(\bar{u})]}{(1-\delta)[1 + e_p^*\delta p - e_i^*\delta p[1-F(\bar{u})] - \delta F(\bar{u})]}.
$$

We employ the interim and ex ante values for L and M by considering the effort levels in each of the three possible equilibrium patterns for a PE. Note, for example, that in the case of no effort in equilibrium (hence, all projects fail) we must have all payoffs equal to zero since all of the value functions collapse to zero when $e_i^* = e_p^* = 0$.

PE **with Full Effort** Consider a PE in which M chooses effort in both the initial state and the persistence state. Substituting $e_i^* = e_p^* = 1$ into the value function for L, we find

$$v^L(u) = \begin{cases} pu + \delta V^L & \text{for } u < \bar{u} \\ \frac{pu+\delta p V^L}{1-\delta(1-p)} & \text{for } u \geq \bar{u} \end{cases}$$

$$V^L = \frac{p\mu[1-\delta(1-p)] + \delta p(1-p)\bar{\mu}[1-F(\bar{u})]}{(1-\delta)[1-\delta(1-p)F(\bar{u})]}$$

Incentives for persistence by L then follow a comparison that is analogous to the full information analysis but with the above payoffs. $IC - L$ requires persistence for projects with u above $\bar{u}$ and termination for those below. As before, this leads to

$$\frac{p\bar{u}}{1-\delta} = V^L \tag{55}$$

and this will determine the threshold for a full-effort PE. Substituting for V^L, we find that (55) reduces to

$$\bar{u}[1-\delta(1-p)F(\bar{u})] - \delta(1-p)\int_{\bar{u}}^{\infty} uf(u)du = [1-\delta(1-p)]\mu.$$

Thus, we find the leader incentives determine a unique threshold for a full-effort equilibrium.

Lemma 14. *There is a unique threshold $\bar{u}_{FE}^*$ that satisfies $IC - L$ (55) in a PE with full effort and this threshold satisfies $\mu < \bar{u}_{FE}^*$.*

Because L is informed about the project value u when deciding whether to persist, $IC{-}L$ binds from above and below at the $\bar{u}$ threshold. In contrast, as we will see, incentives for M are one-sided with respect to effort. Further, effort from M in both the initial and persistence states creates the incentive for L to persist with a threshold strictly above the mean.

Equilibrium incentives for M must support the effort choice in each of the two states. In the persistence state, no effort means failure is certain and, consequently, L will persist with the current project. Thus, $IC_p - M$ only requires $v^M \geq 0 + \delta v^M$, which reduces to $v^M \geq 0$. Referring back to our v^M formula above and substituting with $e_p^* = 1$ we have

$$0 \leq v^M = e_p^* \frac{p\bar{\mu} - c + \delta p V^M}{1 - \delta(1 - p e_p^*)} = \frac{p\bar{\mu} - c + \delta p V^M}{1 - \delta(1 - p)} \qquad \Leftrightarrow \qquad 0 \leq p\bar{\mu} - c + \delta p V^M$$

for $IC_p - M$. Now, substituting with V^M from above with $e_i^* = e_p^* = 1$ and simplifying, we have

$$p\bar{\mu} - c + \delta p V^M \geq 0 \qquad \Leftrightarrow$$

$$p\bar{\mu} - c \geq \frac{-\delta p}{1 - \delta p - \delta(1-p)F(\bar{u})}(p\mu - c)$$

for the $IC_p - M$ incentive constraint. Note that this will be a one-sided constraint on the conditional mean $\bar{\mu}$ and, hence, on the threshold $\bar{u}$ on which $\bar{\mu}$ is based. The nature of the constraint, however, depends on whether the project exceeds or falls below the ex ante static threshold. When $p\mu - c > 0$, so that M would provide effort in a static setting for a project with a random draw from F, then $IC_p - M$ is relaxed and can be satisfied with a relatively low threshold, meaning one that allows for $\bar{\mu}$ that falls below c/p. Intuitively, success implies a new draw, which is an attractive option when the mean is above c/p, and M is willing to provide effort and accept a negative flow payoff on the current project to generate a success. In contrast, when $p\mu - c < 0$, as when the mean is low or cost is high, a new draw is unattractive ex ante. As effort generates a chance of success and, hence, a new project draw, it is the current project that must provide the incentive. As required by the negative term $-\delta p$ in the numerator of $IC_p - M$, we must now have $p\bar{\mu} - c$ above a positive lower bound to provide the work incentive. The equilibrium channel for generating this relationship is via a higher $\bar{u}$ threshold and, in turn, a higher conditional mean $\bar{\mu}$.

Incentives for M in the initial state highlight the uncertainty for M regarding the current project value and whether L will persist or terminate. Shirking initially implies a flow payoff of zero and, in contrast to the certainty in the persistence state, a subsequent chance $F(\bar{u})$ of a new draw for payoff δV^M and $1 - F(\bar{u})$ of persisting for payoff δv^M. Thus, the effort incentive constraint in the initial state, $IC_i - M$, requires

$$V^M \geq 0 + \delta V^M F(\bar{u}) + \delta v^M [1 - F(\bar{u})] \qquad \Leftrightarrow$$

$$V^M \{[1 - \delta F(\bar{u})][1 - \delta(1-p)] - \delta^2 p [1 - F(\bar{u})]\} \geq \delta [1 - F(\bar{u})](p\bar{\mu} - c),$$

upon substituting with the v^M expression. Then, upon substituting with the V^M expression, simplifying, and collecting terms, the $IC_i - M$ constraint reduces to

$$p\mu - c \geq \left[\frac{\delta p[1-F(\bar{u})]}{1+\delta p - \delta F(\bar{u})}\right](p\bar{\mu} - c).$$

Initial state incentives for M contrast strongly with those in the persistence state. First, recall that $IC_p - M$ allowed for a tradeoff between the initial and persistence states in that a negative ex ante initial flow payoff, $p\mu - c < 0$, could be offset by a sufficiently positive persistence payoff flow, $p\bar{\mu} - c > 0$, by having a high cutoff and conditional mean. This will not work for initial state incentives. In $IC_i - M$, a negative ex ante flow payoff implies the same for the persistence payoff flow. That, of course, leads to a contradiction since both payoff flows cannot be negative in equilibrium, as M would never choose to provide effort (formally, $\bar{\mu} > \mu > 0$ always holds so we know the conditional flow payoff must exceed the ex ante flow payoff, $p\bar{\mu} - c > p\mu - c$; but the term in brackets in $IC_i - M$ is between 0 and 1, so $p\mu - c$ must be less negative and hence greater than $p\bar{\mu} - c$). Thus, $IC_i - M$ requires a positive initial flow payoff, $p\mu - c > 0$, and, since the conditional flow payoff is necessarily larger, as $\bar{\mu} > \mu$, initial state effort incentives reduce to the requirement that the initial positive ex ante payoff flow not be too small a fraction of the larger positive conditional flow. Intuitively, the cutoff level is the equilibrium channel for ensuring that this holds.

A comparison of $IC_i - M$ and $IC_p - M$ now reveals that initial incentives for effort are sufficient for persistence incentives. We know that $IC_i - M$ implies $p\mu - c > 0$ so that a necessary condition for a PE with full effort is that the mean project value exceeds the static threshold for effort. Since the conditional mean must exceed the mean, so that $p\bar{\mu} - c > 0$ must hold, we see that $IC_p - M$ is then necessarily satisfied. The intuition, as described above, is that persistence incentives are very relaxed, to the point of permitting a negative persistence payoff flow, once we know that a project is attractive ex ante. Thus, $IC_i - M$ will be sufficient to ensure M effort incentives are satisfied.

With these elements in place, we are ready to characterize PE with full effort.

Proposition 22. *There exists a unique cost level c^* where $0 < c^* < p\mu$ such that if $c \leq c^*$ then a unique full-effort PE exists and $\bar{u}^*_{FE}$ is the cutoff level, where $\mu < \bar{u}^*_{FE}$. If $c > c^*$ then a full effort PE does not exist.*

The static ex ante benchmark for effort, $p\mu - c > 0$, provided a necessary condition for the full effort PE but we now see that this is not sufficient. Instead, an equilibrium with full effort requires a cost $c < c^* < p\mu$ and so projects must meet the more stringent ex ante profitability threshold of $p\mu - c^* > p\mu - c > 0$. This reflects the incentive constraint for initial effort by M. Intuitively, a manager always has the option of not working initially (the "wait" part of this PE). This implies certain failure today but if the project exceeds the threshold, then the leader will persist. If not, then it will be terminated for a new project draw. Effort today shifts likelihood to the success outcome and, hence, a new project draw. But the persistence flow payoff necessarily exceeds the ex ante flow payoff since $\bar{\mu} > \mu$. If $p\mu - c$ is positive but too small, the cost of effort dominates and the manager will shirk, given the $\bar{u}^*_{FE}$ cutoff.

“Wait and See” PE with Initial Shirking Consider a “Wait and See” PE in which the manager shirks initially and provides effort only after persistence by the leader. Substituting $e_i^* = 0$ and $e_p^* = 1$ into the value functions for L, we find

$$v^L(u) = \begin{cases} \delta V^L & \text{for } u < \bar{u} \\ \frac{pu+\delta p V^L}{1-\delta(1-p)} & \text{for } u \geq \bar{u} \end{cases}$$

$$V^L = \frac{\delta p \bar{\mu}[1-F(\bar{u})]}{(1-\delta)[1+\delta p - \delta F(\bar{u})]}.$$

As before, equilibrium incentives require that L persist after a failure when u is above $\bar{u}$, leading to $v^L(u) \geq V^L$. And L must terminate when when u is below $\bar{u}$, leading to $V^L \geq \delta pu + \delta^2 v^L(u)$, since M would provide effort after observing persistence by L. Combining, we find once again that $IC-L$ requires $p\bar{u}/(1-\delta) = V^L$. Substituting for V^L, we find that the project threshold must satisfy

$$\delta \int_{\bar{u}}^{\infty} u f(u) du = \bar{u}[1+\delta p - \delta F(\bar{u})].$$

Lemma 15. *There is a unique threshold $\bar{u}_{SW}^*$ that satisfies $IC-L$ in a PE with initial shirking and this threshold satisfies $\bar{u}_{SW}^* < \mu/p$.*

Thus, leader incentives determine a unique threshold for the equilibrium. Next, given the project threshold of $\bar{u}_{SW}^*$, we must examine when the manager will find it optimal to shirk initially and provide effort after the leader persists with the project.

Incentives for M in the persistence state follow a similar logic to that for the full-effort

PE. No effort implies that failure is certain and L will persist with the current project. Thus, $IC_p - M$ becomes $v^M \geq 0 + \delta v^M$ and reduces to $v^M \geq 0$. Employing our v^M formula from above with $e_p^* = 1$ we again find $p\bar{\mu} - c + \delta p V^M \geq 0$ for $IC_p - M$. Now, however, substituting with V^M from above but with $e_i^* = 0$ and $e_p^* = 1$ and simplifying, we have

$$p\bar{\mu} - c + \delta p V^M \geq 0 \qquad \Leftrightarrow$$

$$p\bar{\mu} - c + \frac{\delta^2 p[1 - F(\bar{u})]}{(1-\delta)[1 + \delta p - \delta F(\bar{u})]}(p\bar{\mu} - c) \geq 0 \qquad \Leftrightarrow$$

$$p\bar{\mu} - c \geq 0$$

for the $IC_p - M$ incentive constraint. This contrasts sharply with M incentives in the persistence state for a full effort PE. In that case, we found a subtle tradeoff between the flow payoff for the project in the initial state and that in the persistence state, where the tradeoff was governed by the threshold and the link to updated beliefs for M. In a PE with initial shirking, however, the initial stage always has a flow payoff of zero. Thus, the $IC_p - M$ incentive constraint reduces to the simple requirement that the project have a positive flow payoff once beliefs have been updated. In turn, the threshold $\bar{u}_{SW}^*$ must be sufficiently large that the conditional mean $\bar{\mu}$ is above c/p.

Turning to the equilibrium incentive constraint to shirk initially, $IC_i - M$, we compare V^M to the deviation payoff for M from providing effort:

$$V^M \geq p\mu - c + \delta p V^M + \delta(1-p)[F(\bar{u})V^M + [1 - F(\bar{u})]v^M] \qquad \Leftrightarrow$$

$$p\bar{\mu}-c \geq \left[\frac{1+\delta p-\delta F(\bar{u})}{\delta p[1-F(\bar{u})]}\right](p\mu-c)$$

where we have substituted with $e_i^* = 0$ and $e_p^* = 1$ for v^M and V^M from above and simplified. The term in brackets on the right always exceeds 1, so $IC_i - M$ requires that the flow payoff for the project in the persistence state, after beliefs have been updated, exceed a multiple of the flow payoff to effort in the initial state. In fact, we see that the inequality has simply been reversed from the $IC_i - M$ constraint for the full effort PE. There, we found that the $IC_i - M$ constraint was sufficient for the $IC_p - M$ so that a manager who was willing to provide initial effort would necessarily find it optimal to provide effort in the persistence state. Incentives in a PE with initial shirking are more complex.

Collecting the two incentive conditions for reference, we have

$$p\bar{\mu}-c \geq 0 \qquad\qquad IC_p - M, \tag{56}$$

$$p\bar{\mu}-c \geq \left[\frac{1+\delta p-\delta F(\bar{u})}{\delta p[1-F(\bar{u})]}\right](p\mu-c) \qquad\qquad IC_i - M. \tag{57}$$

This highlights that there will be two incentive scenarios for a PE with initial shirking. For the first one, suppose that a project fails to meet the ex ante static threshold, $p\mu - c < 0$. In this case, $IC_p - M$ will be sufficient. With $p\bar{\mu} - c \geq 0$, so that the project meets the static threshold but with the updated belief, $IC_p - M$ is satisfied and the manager will provide effort once the leader persists with the project. In turn, $p\bar{\mu} - c \geq 0$ implies that $IC_i - M$ is satisfied since $p\mu - c < 0$. For the second, suppose instead that $p\mu - c > 0$ and the project meets the static threshold. Now, $IC_p - M$ is redundant once initial incentives are satisfied since $\bar{\mu}$ always exceeds μ. But with $p\mu - c > 0$, there is an obvious temptation for the manager to work initially. After all, the leader will persist if u is high and terminate if u is low and the manager does not observe u. What deters initial effort? The answer is that initial effort runs the risk of replacing the current project with a new draw. This would occur whether u is high or low. Shirking, however, ensures an initial failure and the persistence decision follows. As a result, the leader will terminate only the low value projects and persist with the high value ones. This project selection, as a result of shirking, is how the manager benefits. As $IC_i - M$ reveals, when $p\bar{\mu} - c$ sufficiently exceeds $p\mu - c$, the selection benefit makes shirking optimal.

We can now characterize PE with initial shirking.

Proposition 23. *There exist cost thresholds c^* and c^{**}, where $0 < c^* < p\mu < c^{**}$, such that if $c^* \leq c \leq c^{**}$ then there exists a unique PE with initial shirking and $\bar{u}^*_{SW}$ is the cutoff level. If $c < c^*$ or if $c^{**} < c$, then a PE with initial shirking does not exist.*

A PE with shirking thus arises in settings where the cost c is not too extreme on either the low or the high side. When cost is very low, the incentive for initial effort is strong and the PE with full effort exists. As we see later, when cost is very high, incentives collapse and the PE with no effort exists. Equilibrium shirking arises between these extremes and, intriguingly, it applies for a range of c below and above the static benchmark at $p\mu$. Thus, when c is above c^* but below $p\mu$, equilibrium will have shirking in the initial stage even though the static ex ante payoff to effort is positive. When c is higher, above $p\mu$ but below c^{**}, initial shirking is consistent with the ex ante threshold. As discussed above, these patterns reflect the two incentive scenarios and which of the initial state and persistence state incentive constraints for M is binding when the project threshold is $\bar{u}^*_{SW}$.

A second point of contrast relative to the full effort PE concerns the project threshold. We know that $\mu < \bar{u}^*_{FE}$ so that the full-effort PE always has a threshold that is above the mean and, hence, the static $p\mu$ benchmark. For the PE with initial shirking, we have

Corollary 5. *The project threshold for PE with initial shirking, $\bar{u}^*_{SW}$, satisfies*

(i) if $\delta > \frac{p}{1-p^2}$ then $\bar{u}^*_{SW} > p\mu$;

(ii) if $\delta < \frac{p}{1+p-p^2}$ then $\bar{u}^*_{SW} < p\mu$;

(iii) $\bar{u}^*_{SW} < \mu/p$

(iv) $\bar{u}^*_{SW} < \bar{u}^*_{FE}$

.

The point of contrast is that equilibrium shirking can occur with a project threshold that is below $p\mu$. In the full-effort PE, the project threshold always exceeds $p\mu$. From Corollary 5, we see that a sufficiently high chance of project success, p, relative to the discount factor, δ, is sufficient for $\bar{u}^*_{SW}$ to be below $p\mu$. As the threshold is not dependent on c, we can easily assess the range of projects for which the leader will persist, $u \geq \bar{u}^*_{SW}$, in relation to the range of costs, $c^* < c < c^{**}$, for which equilibrium shirking arises. Recalling that $p\mu$ always lies between c^* and c^{**}, we then see that equilibrium shirking will exist in settings where $p\mu$ is strictly below c, namely, when the cost c is close to c^{**}. Thus, we will have $\bar{u}^*_{SW} < p\mu < c$ in this case and the equilibrium path will involve project draws where the value u is between $\bar{u}^*_{SW}$ and $p\mu$, so that $pu - c < 0$, and the manager, after initially shirking, will switch to providing effort after the leader persists.

In the wait and see equilibrium, persistence is met with work. However, persistence (and work) only obtain when the expected value of success from the updated distribution of project values $\bar{\mu} \geq c$. But in our earlier complete information analysis regarding low-$\bar{u}$ (dim prospects) equilibria, persistence thresholds occurred at levels in which the manager's expected value of success from the existing project $pu \leq c$. Why does incomplete information lead to a qualitatively different $\bar{u}$? The reason is that incomplete information changes the relative attractiveness of a future project relative to the current project. Under complete information, the manager has an incentive given persistence to succeed with the current project to gain a more attractive future project. But under incomplete information, the current project conditional on persistence is more attractive in expected value terms than a new draw from the unconditioned distribution with a mean value of μ.

PE with No Effort This is, of course, the simplest case to analyze. With no effort in either state, $IC - L$ is trivially satisfied. Since the choice to persist or not has no impact on effort, L always has a payoff of zero and $IC - L$ holds for any value of $\bar{u}$. For the manager, $IC_p - M$, the incentive to shirk in a persistence state, reduces to $0 \geq p\bar{\mu} - c$ and this does restrict the range of equilibrium cut-off levels. For the initial state, $IC_i - M$ requires shirking and this reduces to $0 \geq p\mu - c$. Note that this is a pure parameter condition and does not involve $\bar{u}$. The characterization for PE with no effort is then

Proposition 24. *If $p\mu > c$ then a PE with no effort does not exist. If $p\mu \leq c$ then there exists a set of no-effort PE with each PE in this set indexed by a cutoff level $\bar{u} \in [0, \bar{u}_0^*]$.*

The upper bound $\bar{u}_0^*$ *is defined by*

$$p\int_{\bar{u}_0^*}^{\infty} uf(u)du = c[1-F(\bar{u}_0^*)]$$

and satisfies $0 < \bar{u}_0^* < c/p$. *All* PE *in this set are payoff-equivalent, with a payoff of zero for* L *and* M, *and involve no effort in either of initial or persistence project states. Finally, there does not exist a no-effort* PE *with a cutoff above* $\bar{u}_0^*$.

The exact details of the set of no-effort equilibria should not obscure the basic simplicity of these equilibria. If projects are attractive to M on a one-shot basis, meaning $p\mu > c$, then initial effort is a profitable deviation for M and a no-effort PE cannot be supported. Thus, a "low" project value mean μ relative to cost c is necessary for a no-effort PE. It is also sufficient, once we recognize that the persistence threshold must not support effort by M based on updated project beliefs. This is why the upper bound $\bar{u}_0^*$ is necessarily below c/p. Any threshold above $\bar{u}_0^*$ will lead to a belief for M, after persistence by L, with a conditional mean that is above c/p and effort would then be optimal for M. At $\bar{u}_0^*$ the conditional mean is exactly c/p, so any smaller threshold makes shirking the optimal choice for M.

The existence of a no-effort PE when $p\mu < c$ under incomplete information contrasts with equilibrium outcomes under full information. When M can base the effort choice on the observed u, a sufficiently low draw always leads to no effort in equiibrium. But termination then leads to a new project draw and positive payoffs for L and M as effort will be forthcoming on attractive future projects. In contrast, the no-effort PE has zero payoffs for both of L and M. This collapse of effort is a distortion introduced by incomplete information. The effort choice must be based on an expectation rather than the actual project draw u. The initial mean and then the updated mean both support shirking by M, and given this effort pattern, L is indifferent regarding persistence since it does not influence effort incentives.

Thus, implicit coordination again emerges as an equilibrium feature when future prospects are "dim" in the sense that the mean project falls below a threshold. In the complete information setting, with Proposition 21, we found an equilibrium set of thresholds. In that setting, where the project value was known to L and M, the choice to work or shirk by M could be tailored according to whether u was above or below $\bar{u}$. Coordination was reflected in the range of thresholds, $\bar{u} \in (\bar{u}_0, c/p)$, that could support equilibrium behavior. Intuitively, with a common reference of $\bar{u}$, L would shirk when u was below $\bar{u}$ because L was expected to terminate. Above $\bar{u}$, M provided effort and L would persist. Because any such threshold was below c/p, M necessarily shirks below the threshold (in contrast to the "bright future" case where M worked for a range below the threshold). Under incomplete information, M must choose based on an expectation of the project value. Implicit coordination is still preesent in equilibrium but the equilibrium threshold range is distorted down. Now, the threshold must be between $\bar{u} \in [0, \bar{u}_0^*]$ and the upper bound is strictly smaller than c/p. Coordination is reflected in the reinforcement of the persistence and effort decisions. Since M never provides effort, L is indifferent across thresholds. M, however, is sensitive to the updated belief. Since the equilibrium range of thresholds lies near 0, as long as L persists and terminates according to the threshold, the updated belief of M remains sufficiently small that shirking is optimal. Of course, as noted above, the main distortion relative to complete information is that coordination is now concentrated on the no-effort outcome in equilibrium.

4.6 Persistence with Team Production

We now consider persistence in a setting with team production. Suppose that the R&D unit now consists of two managers, 1 and 2, who independently choose whether or not to exert effort on projects that L brings forward. Our goal is to study, in isolation, the effects of increased agency within the organization on equilibrium persistence. Therefore, it is helpful to suppose that each manager's effort contributes probability $p/2$ to the likelihood of success on a given project and that the cost of exerting effort is $c/2$ for each manager. In other words, suppose that total productivity and costs of the R&D unit (given full effort) is identical to the baseline model with one manager. The only difference is that achieving success with probability p requires that two managers exert effort rather than just one.

A convenient feature of our team production setting is that each manager's static incentive to work (i.e., when L terminates) is identical to the one-manager case. Indeed, if L is expected to terminate a project with value u, then the best manager i can do is maximize current-period expected utility, namely $(p/2)(e_i + e_j)u - (c/2)e_i$. From this, we conclude that effort is optimal if and only if

$$\frac{p}{2}u - \frac{c}{2} \geq 0 \iff u \geq c/p.$$

Dynamic incentives to work when L persists, by contrast, are lower than in the one-manager case due to the ability of manager i to free-ride on the efforts of manager j. Indeed, if L is expected to persist on a project with value u, then manager i's best response is to maximize interim expected utility, namely

$$v_i(u) = \frac{\frac{p}{2}(e_i + e_j)(u + \delta V_i)}{1 - \delta(1 - \frac{p}{2}(e_i + e_j))}.$$

Maximizing with respect to e_i reveals that $e_i = 1$ is (dynamically) optimal if and only if

$$u \geq \frac{c}{p} - \delta V_i + \frac{\delta}{1-\delta}\frac{c}{2}e_i. \tag{58}$$

Manager i's incentive to work under persistence is determined by three terms. The first term, c/p, naturally captures manager i's static incentives. The second term, $-\delta V_i$, captures the additional incentive to work in order to accelerate the arrival of new projects. The third term, $\frac{\delta}{1-\delta}\frac{c}{2}e_j$, captures manager i's incentive to free-ride on the efforts of manager j. Intuitively, free-riding incentives are largest whenever managers are relatively patient (i.e. δ is close to 1) and effort is relatively expensive (i.e. c is large).

We are now in a position to characterize equilibrium persistence in the two-manager setting. As before, we distinguish between two cases: when the future is “bright" and when the future is “dim."

Proposition 25. *Suppose that*

$$\frac{c}{p}\left[1-\delta(1-p)F\left(\frac{c}{p}\right)\right] < \int_{c/p}^{\infty} uf(u)du \tag{59}$$

holds. Then there exists a unique "full-effort" PE. In this equilibrium, each manager works iff $u > c/p$ *and* L *persists iff* $u > \bar{u}$*, where* $\bar{u} \in (c/p, \infty)$ *is identical to the one-manager case.*

When the future is bright, Proposition 25 shows that each manager's incentive to accelerate the arrival of future (high expected value) projects always dominates the incentive to free-ride. Thus, the only binding equilibrium constraint in this case is the leader's IC condition, $p\bar{u} = (1-\delta)V_L$. Importantly, the solution to this IC condition is identical to the one-manager setting since both managers exert effort using the statically-optimal effort rule, $e = 1 \iff u > c/p$.

An immediate corollary of Proposition 25 is that a dim future is necessary for free-riding to occur in PE. Our next result shows that a dim future is also sufficient.

Proposition 26. *Suppose that*

$$\frac{c}{p}\left[1-\delta(1-p)F\left(\frac{c}{p}\right)\right] > \int_{c/p}^{\infty} uf(u)du \tag{60}$$

holds. Then free-riding occurs with positive probability in every PE. In particular, every pure-strategy PE features two thresholds, $\bar{u}$ *and* u^**, where* $\bar{u}^* \in (\bar{u}, \infty)$*, such that (i)* L *persists iff* $u > \bar{u}$*, (ii) both managers work iff* $u > (c/p, \min\{c/p, \bar{u}\}) \cup (u^*, \infty)$*, and (iii)* *only one manager works iff* $u \in (\bar{u}, u^*)$.

When the future is dim, the motivational effects of persistence are relatively weak compared to free-riding incentives because any new project is expected to have low value. For this same reason, L is willing to persist on projects with value far below c/p. These two effects combine to guarantee that free-riding occurs with positive probability in every possible PE. More precisely, there always exists an interval of low-value projects on which L is willing to persist on a project that she knows receives only partial effort.

4.7 Discussion

In this section we offer an interpretation of the persistence equilibrium through the lens of commitment and strategy. We then discuss some factors that relate to R&D project termination decisions that were suppressed in the model and how they might impact the intuition and results.

Endogenous Commitment and Strategy One challenge facing a leader is getting the rest of the organization to follow the leader's direction. If persistence is a valuable instrument for generating effort on a project-by-project basis, is persistence also valuable for implementing strategy?

Ghemawat (1991) argues commitment is the defining feature of strategic decisions because true commitments shape the path of future firm decisions. He gives as examples physical and economic commitments (e.g. building production capacity or engaging in long-term contracts) which serve this commitment purpose if they are not easily reversed or repurposed to another project. But irreversibility is not a necessary condition for creating the reliability critical for strategy (Van den Steen 2017) and persistence, in equilibrium, can serve a commitment role.

In the R&D setting analyzed above, the decision to persist with an unsuccessful project is sufficient to persuade the subordinate that the leader will persist with the project going forward until it becomes successful. This equilibrium outcome gives the leader traction to induce more effort from the subordinate than the subordinate would deliver given short-term incentives. Persistence becomes a valuable policy instrument to expand the range of projects to which a leader and, hence the firm, is "committed."

The incomplete information "wait and see" equilibrium is instructive in understanding this endogenous mechanism. In this equilibrium a manager only exerts effort in response to a persistence choice by the leader. Until persistence is observed, the manager is rightfully skeptical about whether the leader is committed to the project at hand because some (low value) projects will be abandoned by the leader in the next period. But subsequent persistence signals that the project has a high enough value for the leader to persist until success is achieved—the leader is committed in equilibrium to the project—so a manager will exert effort on this project going forward.

This incomplete information analysis captures a legitimate skepticism that a manager might have regarding an assigned project when the value the leader has for the project is quite uncertain to the manager. In practice, other factors may reinforce the impact of uncertainty which we do not model. For example, the leader (or the organization) may have a history of overhyping its projects or abandoning attractive, but long-term, projects when short-term performance pressures are high. On the other hand, in settings in which an initial project (as opposed to what one might expect in a series of initial projects) is much ballyhooed or thought to be particularly important to the leader, there may be an initial burst of effort that might offset the "wait and see" incentive of the manager.

Context for Project Development and Termination We now discuss several features of the R&D project development and termination processes that were simplified or suppressed in our model. These include compensation-oriented incentives, binary treatment of effort, behavioral biases towards persistence, and a more complex intra-organizational decision-making environment.

Compensation and Payoff Division Our focus on the impact of the persistence instrument led us to suppress consideration of optimizing incentive design. In our model, payoffs to successful projects are specified as a deriving from an exogenous split of successful project payoffs between the leader and the subordinate. It would be easy to parameterize the actual split rather than treat it as equal as we do in the base model, but we do not consider how a different payoff structure beyond a simple split might better optimize effort. We also do not consider potential issues associated with implementing the ex post split—essentially, we appeal to the idea that there is a promise to pay the subordinate the split, perhaps via a relational contract (Baker, Gibbons, and Murphy 2002).

The purpose of our analysis is to explore the impact of persistence, not that of compensation design. The management literature (e.g. Bower 1970) makes clear that the project-by-project persistence decision of the leader (principal) is a first-order policy decision. Our analysis shows that persistence impacts subordinate (agent) effort. Hence, we see our work as complementing the work on compensation design (Manso 2011, Hellmann and Thiele 2011) which focuses on commitments that alter the level or type of agent effort. The integration of the two instruments—persistence and incentive design—is left for future work.

Treatment of Effort The interpretation of c, the cost of effort, by the manager deserves some discussion. The cost of effort can be thought of as the manager's private opportunity cost of effort. This interpretation encompasses the value of leisure, but we prefer to think of it as the cost of working on the focal project versus working on a different project that is not directly sanctioned by the leader (Burgelman 1991; Bower and Gilbert 2005) or perhaps on an incremental value project that was also assigned to the manager. Many of these unsanctioned projects could have value to the firm, hence, one can interpret c as the difference between the payoff of the project to the leader from a success on the focal project and any value emerging from the unsanctioned project versus the payoff to the manager. For example, the manager responsible for the unexpected success of the unsanctioned project presumably receives relatively more credit for its success than does the leader who has publicly championed the focal project.

It is convenient in our model to assume a binary choice of effort. What matters for the results is having different levels of effort rather than have full or no effort. The actual interpretation of level of effort depends on how the "manager" in the model corresponds to roles in an organization. One interpretation is that the manager is representing the R&D engineers or scientists who do the R&D work. In this case effort could be thought of in terms of the amount of work that is done on a project. These knowledge workers may be driven to work harder on projects that are seen as more important or more pressing (see, e.g., Buell 1967, Balachandra and Raelin 1980), or may be demotivated if they think the project is likely to be soon terminated or, worse yet, that they might be soon terminated. In the latter case the employee would likely substitute job search for project work.[19]

Our preferred interpretation is that the manager in our model manages more than one project. Then, lack of effort can be interpreted as the assignment of better personnel to other projects, various subterfuges, or straight out action against stated policy.[20] Bower and Gilbert (2005) offer examples of specific directives of the CEO that were deflected by his managers (e.g. the CEO of Teradyne, a semiconductor testing equipment manufacturer, could not get his division managers to put their best talent on CEO-favored projects involving cheaper but inferior test equipment), while Burgelman (1991) found that managers at Intel disguised work on RISC chip as development of a coprocessor to circumvent Intel's strategy not to pursue RISC based businesses.[21] Gilbert (2006) also describes how a

---

[19] See, e.g., Sarah Perez, "Google cancels half the projects at its internal R&D group Area 120," Techcrunch.com, September 14, 2022, Kate Conger, Ryan Mac, and Lauren Hirsch, "Twitter Tries Calming Employees as Deal with Elon Musk Looms," New York Times, October 21, 2022.)

[20] This inability to force work only on a narrowly assigned project not only reflects lack of observability, it may also reflect a lack of clarity about the specifics of a manager's assignment and poor communications, as well as conflicting spheres of influence (Bower and Gilbert 2005) or project teams that have been given somewhat conflicting priorities (Verma 2011).

[21] Other famous examples include self-development of blue light by inventor at Hitachi and laser printer at Xerox NY research center.

newspaper publisher's online project was unable to get needed help from the newspaper operations people because they thought the core business was much more important.

General Behavioral Biases toward Persistence We model the decision makers and employees as risk-neutral payoff maximizers whose decision making does not suffer from behavioral biases. We believe this is a good foundation on which to develop an understanding of the role and use of persistence as a management policy instrument. In principle, one could extend the model by introducing one or more of the behavioral biases (e.g., Staw and Ross 1980, Guler 2007, 2018) that and others have identified regarding project termination and see how they affect the outcomes. One finding from the behavioral decision bias literature as applied to our context suggests that organizations are biased against termination—there is an escalation of commitment over time (see, e.g., Sleesman, et al. 2017 for a review). This bias seems to operate at all levels, affecting team members as well as corporate leaders who have termination or resource allocation authority. Some of these biases may have an underlying individual rational basis (e.g., team members suffer a cost when their projects are terminated, projects are important for gaining power within the organization) but are dysfunctional at the organization level. Other biases stem directly from individual decision-making defects (e.g., inertia, not revisiting previous decisions, cognitive dissonance, and the like).

If the net effect of these biases is to push the organization to terminate existing projects less often, then decision makers, from the perspective of our model, persist more often than they should. But while the biases may affect the absolute level of persistence, persistence as a policy instrument is not obviously affected. Consider how the bias might operate in the context of our model. The bias can be seen as altering current versus future project payoffs: decision makers incorrectly perceive the payoffs from the current projects to be relatively greater than its actual value. Projects will still be terminated by the leader, but effectively at a lower perceived value cutoff level. From the viewpoint of the subordinate, there will still be projects for which the cost of effort is greater than the perceived value, so the basic forces analyzed above are still present.[22]

Organizational Complexity, Agency, and Information Structure The agency problem faced by firms is evident in many detailed descriptions of R&D and project histories in the literature. In the strategy literature, for example, there is an extensive literature on emergent strategies (Mintzberg and Waters 1985), some versions of which recognize implementation problems that interfere with top-down strategy directives relating to grand strategy but also to specific implementations of that strategy as represented in various corporate programs or projects. A theme in these studies is that upper management has blunt rather than targeted instruments for achieving its ends. Our model reflects this perspective by adopting a moral hazard modeling approach in which subordinate effort cannot be observed.

The resource allocation process (RAP) of Bower (1970) is a description of how strategy

---

[22] One bias identified in the literature involves the cost incurred by project team members when their project is exited through failure rather than through success (Peck et. al. 2015). In the complete information model, the subordinate anticipates if the leader will exit and will have an additional incentive to work given the added value of exiting through success rather than through failure. In the incomplete information variant of the model, one can think of this bias as lowering the expected cost of effort because if no effort is expended, failure is assured so whenever the leader chooses not to persist with the current project (which only occurs if the project is unsuccessful), the subordinate incurs a cost independent of effort. The subordinate again has even more incentive to exit via success than failure which can be captured as a reduction in expected cost of effort.

emerges from the iterative actions of managers at multiple levels of a (large) organization. This description draws from numerous case studies of strategy development and change in large organizations. Power and information are dispersed across these managers. Strategy is seen as the result of the allocation of resources to various projects and initiatives. Because top leaders lack detail regarding the projects under consideration for these resources, they rely heavily on the recommendations—and reputations—of the lower-level general managers. RAP scholars discuss various examples of firms where commitments made by general managers lower in the organization—often against current strategy–led to changes to the firm's official strategy. These commitments were made in part because "decision rights are almost inevitably in conflict" (Bower and Gilbert 2005 p18) and "commitments must be made. (p13)"

In the revised RAP model, Bower and Gilbert (2005) focus on two core processes: definition and selection. Definition concerns the "content of strategic thinking" (p.442) which can be thought of as a plan to solve some perceived problem faced by the firm and consists of aligning, translating, and initiating activities. Selection is "focused on the choice of projects and business plans to propose to sponsor and to approve funding" (p 443) which includes championing a good idea, getting support from general managers which can be seen as the company's "merchant bankers", and obtaining "at least formal corporate approval." (p 447)

In RAP terms, persistence in our model is analogous to a firm's continued resource commitment to a particular project which Bower and others in the RAP tradition see as the defining decision regarding strategy choice. Our “leader” is equivalent to the mid-level general manager who is the focus of the selection part of the process and the “manager” would be the lower-level manager who has direct project authority if given resources but may also have a portfolio of projects. Our model does not capture the direct competition among various possible projects for funding or the reliance on subordinate manager reputation in the assessment process because we model only one manager. Rather one can think of our model as one where the “next draw from the distribution of projects” comes from the manager or is offered by an advanced project manager who leads a different unit than the implementing manager.

Bower and Gilbert discuss how a firm’s realized strategy emerges from an iteration between definition and selection. Our model operates at a level below that of a firm’s full strategy so does not address this part of the RAP model.

The description of the RAP process suggests that the subordinate has better information regarding the details of the project than the leader (mid-level general manager). What matters for our persistence model is less the details than the expected valuation of the project. In the RAP model the mid-level general manager relies on the subordinate’s reputation which is one way in which the subordinate’s assessment of the project to the mid-level manager could be viewed as reasonably credible. If so, then the project level information held by the subordinate can reasonably be assumed to be not private (it is accurately presented) which is how we model the subordinate’s information. On the other hand, RAP acknowledges that upper-level managers have other information from other parts of the firm that is not available to their subordinates. Hence, ultimate value of a project (and certainly the leader's perceived value of the project) can be viewed as private information of the leader that is material to the choices of the subordinate. More specifically, observers of the R&D management process note the importance of basing decisions about projects on information that spans several projects (Jekunen 2014 or where a single project is a "building block" for multiple projects (Brennan, et. al. 2020). Thus, while the technical feasibility of a project may be best known to the immediate project team, the overall value of the project is better known to their superiors. This view is also reflected in the key criteria (e.g. strategic fit) firms use to evaluate their projects (Cooper, Edgett, and Kleinschmidt 1998).

4.8 Appendix: Proofs and Derivations

Proof of Lemma 11. Consider a PE with threshold $\bar{u} \geq 0$ and suppose the current project has payoff $u \leq \bar{u}$. Because L will terminate this project, there will be a new draw next period whether or not M works and, in turn, whether or not the project is a success. Thus, the only effect of effort on the payoff of M is with the current project. This implies M will work if and only if $u > c/p$. The implication is immediate: in a high-$\bar{u}$ equilibrium, where $\bar{u} > c/p$, the equilibrium must have $e^*(u_t) = 1$ for $u_t \geq c/p$ and $e^*(u_t) = 0$ for $u_t < c/p$, and the claim in Lemma 11 is established. □

Proof of Proposition 20. We know from the text that $e^*(u)$ as specified in Lemma 11 for a given $\bar{u} > c/p$ implies that incentive compatibility for M is satisfied. Also, from the text, to satisfy incentive compatibility for L, the threshold $\bar{u}$ must satisfy (49). Simplifying with our formula for $V^L(\bar{u})$ from the text, we must have the equilibrium threshold $\bar{u}$ such that

$$\bar{u}[1-\delta(1-p)F(\bar{u})]+\delta(1-p)\int_{c/p}^{\bar{u}} uf(u)du=\int_{c/p}^{\infty} uf(u)du.$$

Note that the right-hand-side (RHS) is a strictly positive constant with respect to the threshold. Consider then how the LHS varies with $\bar{u}$. Calculating, we find that the LHS is strictly increasing in $\bar{u}$ with partial derivative equal to $1-\delta(1-p)F(\bar{u})>0$ from a value of $\frac{c}{p}\left[1-\delta(1-p)F(\frac{c}{p})\right]$ for the LHS at $\bar{u}=c/p$ to a limit of ∞ as $\bar{u}\to\infty$. Consequently, there is no such $\bar{u}$ equilibrium threshold above c/p if condition (50) fails and $\int_{c/p}^{\infty} uf(u)du$ is smaller than the value of the LHS at $\bar{u}=c/p$. By monotonicity of the LHS, when (50) holds, there is a unique $\bar{u}>c/p$ that solves the condition for the equilibrium threshold. □

Proof of Corollary 3. We need only show that $\frac{c}{p}[2-\delta(1-p)]<\mu$ implies (50). Note first that

$$\frac{c}{p}[1+F(u)-\delta(1-p)F(u)]<\frac{c}{p}[2-\delta(1-p)]$$

since the LHS is increasing in u and the the RHS is the limit of the LHS as $u\to\infty$. Hence, we have, taking $u=c/p$,

$$\frac{c}{p}\left[1+F(\frac{c}{p})-\delta(1-p)F(\frac{c}{p})\right]<\mu=\int_{0}^{c/p} uf(u)du+\int_{c/p}^{\infty} uf(u)du \qquad \Rightarrow$$
$$\frac{c}{p}\left[1-\delta(1-p)F(\frac{c}{p})\right]<\int_{c/p}^{\infty} uf(u)du-\left[\frac{c}{p}F(\frac{c}{p})-\int_{0}^{c/p} uf(u)du\right]$$

and we are done as the last bracketed term is positive, leaving (50) as the implication. □

Proof of Lemma 12. Suppose $u < \bar{u}$. Then L will terminate and the current work choice by M has no effect on future payoffs. As a result, M optimally chooses not to work since $pu - c < 0$ when chooses $u < \bar{u} \leq c/p$ and we have $e^*(u) = 0$. The proof for the case of $u \geq \bar{u}$ is also simple. Suppose that $e^*(u) = 0$. Then the project will fail with certainty but L will persist. But then, with $e^*(u) = 0$, this will repeat in the next period and in all subsequent periods, implying a payoff of 0 in this u state for M and for L. But this cannot be an equilibrium. M and L each have a positive ex ante payoff. To see why, note that M always has the option of working in states u' where $u' \geq c/p$. Since the initial $t = 1$ draw is from F, by following this work rule M has an ex ante payoff of at least $\int_{c/p}^{\infty} [pu - c] f(u) du > 0$, which in turn implies that L has a positive ex ante payoff. Thus, it would not be optimal for L to persist in u and we have a contradiction. As a result, we must have $e^*(u) = 1$. $\square$

Proof of Lemma 13. Consider (52) from the text

$$p\bar{u} - c + p\delta\left[\frac{\int_{\bar{u}}^{\infty} puf(u)du}{(1-\delta)[1-\delta(1-p)F(\bar{u})]}\right] \geq 0$$

as a function of $\bar{u}$. First, at $\bar{u} = c/p$, this reduces to

$$p\delta\left[\frac{\int_{c/p}^{\infty}(pu-c)f(u)du}{(1-\delta)[1-\delta(1-p)F(c/p)]}\right]$$

which is strictly positive. By continuity, (52) necessarily holds for an interval below c/p. Next note that the LHS of (52) is strictly increasing in $\bar{u}$ because i) $p\bar{u} - c$ is increasing,

ii) the integral in the numerator is increasing at rate $-(p\bar{u}-c)f(\bar{u}) > 0$ for $\bar{u} < c/p$, and iii) the denominator is decreasing at rate $-\delta(1-p)F(\bar{u})$. The remaining question is then whether the LHS is positive or negative at $\bar{u} = 0$. Simplifying the LHS with $\bar{u} = 0$ and $F(0) = 0$ yields that LHS is positive $\Leftrightarrow$

$$\frac{1}{1-\delta}\left[\delta p^2\mu - c[1-\delta(1-p)]\right] > 0$$

which then reduces to the sufficient condition in the Lemma. □

Proof of Proposition 21. We know from Lemma 13 that IC for M holds over an interval below c/p. If we can show the same for L, an interval of the form $[\bar{u}_L, c/p]$ then we will have a set of low-$\bar{u}$ PE where $u \in [\max\{\bar{u}_M, \bar{u}_L\}, c/p]$. Write (53) as

$$\bar{u}[1-\delta(1-p)F(\bar{u})] - \int_{\bar{u}}^{\infty} u f(u) du \geq 0$$

and note that the LHS above is strictly increasing in $\bar{u}$ at rate

$$1-\delta(1-p)F(\bar{u}) + \bar{u}F(\bar{u})\delta(1-p) > 0.$$

At $\bar{u} = 0$, the LHS of the IC condition reduces to $-\mu$ and we see IC for L can never extend to 0. At $\bar{u} = c/p$, we simplify the LHS of the IC condition to find that IC for L reduces to (54). Together with monotonicity and the value of $-\mu$ at $\bar{u} = 0$, this implies that there exists a unique $\bar{u}_M \in (0, c/p)$ at which IC for L binds, thus establishing the result. □

Proof of Lemma 14. It is straightforward to verify that there is a unique $\bar{u}$ that satisfies IC_L:

$$\bar{u}[1-\delta(1-p)F(\bar{u})]-\delta(1-p)\int_{\bar{u}}^{\infty} uf(u)du = [1-\delta(1-p)]\mu.$$

The right-hand side is a positive constant while the left-hand side is strictly increasing in $\bar{u}$, with slope $1-\delta(1-p)F(\bar{u}) > 0$. At $\bar{u} = 0$, the left-hand side is negative at $-\delta(1-p)\mu < 0$. As $\bar{u} \to \infty$ the first term dominates and the left-hand side also goes to ∞. Thus, the left-hand side will cross, from below, the value $[1-\delta(1-p)]\mu$ exactly once. Let $\bar{u}^*_{FE}$ denote this unique $\bar{u}$ value. To see that $\mu < \bar{u}^*_{FE}$ holds, simply evaluate the left-hand side at $\bar{u} = \mu$ to find that

$$\begin{aligned}\mu[1-\delta(1-p)F(\bar{u})]-\delta(1-p)\int_{\mu}^{\infty} uf(u)du < [1-\delta(1-p)]\mu \qquad &\Leftrightarrow \\ \mu[1-F(\bar{u})] < \int_{\mu}^{\infty} uf(u)du &\end{aligned}$$

and the last inequality is clearly valid (integral lower bound). Thus, the crossing $\bar{u}^*_{FE}$ must be at a $\bar{u}$ above $\bar{u} = \mu$. □

Proof of Proposition 22. We know from Lemma 14 that $IC - L$ is satisfied if and only if the threshold is $\bar{u}^*_{FE}$. From the text, we know that $p\mu - c > 0$ is necessary for existence and then that $IC_i - M$ is sufficient for $IC_p - M$ to hold. Consider, then, when $IC_i - M$ is satisfied

$$p\mu - c \geq \left[\frac{\delta p[1-F(\bar{u})]}{1+\delta p-\delta F(\bar{u})}\right](p\bar{\mu}-c) \qquad \Leftrightarrow$$

$$\mu \geq \frac{c}{p} + \left[\frac{\delta p[1-F(\bar{u})]}{1+\delta p - \delta F(\bar{u})}\right]\left[E(\bar{u}) - \frac{c}{p}\right]$$

where we are defining $\bar{\mu} = E(\bar{u}) = \int_{\bar{u}}^{\infty} uf(u)du/[1-F(\bar{u})]$ to make explicit the dependence of the conditional mean on the cutoff; note that $\mu = E(0)$. It is straightforward to verify that the right-hand side above is increasing in c as $0 < [\delta p[1-F(\bar{u})]]/[1+\delta p - \delta F(\bar{u})] < 1$. At $c = 0$, the value of the right-hand side is below μ :

$$0 + \left[\frac{\delta p[1-F(\bar{u})]}{1+\delta p - \delta F(\bar{u})}\right][E(\bar{u}) - 0] = \left[\frac{\delta p}{1+\delta p - \delta F(\bar{u})}\right]\int_{\bar{u}}^{\infty} uf(u)du < \mu$$

since $\delta p < 1 + \delta p - \delta F(\bar{u})$ and $\bar{u} \geq 0$. At $c = p\mu$, we have

$$\begin{aligned}
&\frac{p\mu}{p} + \left[\frac{\delta p[1-F(\bar{u})]}{1+\delta p - \delta F(\bar{u})}\right]\left[E(\bar{u}) - \frac{p\mu}{p}\right] \\
&= \mu + \left[\frac{\delta p[1-F(\bar{u})]}{1+\delta p - \delta F(\bar{u})}\right][E(\bar{u}) - \mu] \\
&> \mu
\end{aligned}$$

and the right-hand side is above μ. hence, there is a cross of μ, from below, as c varies by the value of the right-hand side at a unique c between 0 and $p\mu$; let c^* denote the crossing c value. Note that c^* depends on $\bar{u}$ and we will use $\bar{u} = \bar{u}^*_{FE}$ for the PE.

To complete the proof, let $c < c^*$ and consider whether $\bar{u}^*_{FE}$ as the threshold and $e^*_i = e^*_p = 1$ for effort constitute a PE. We know $IC - L$ is satisfied if and only if $\bar{u} = \bar{u}^*_{FE}$. Since $c < c^* < p\mu$, we know $p\mu - c > 0$ and $IC_p - M$ is satisfied. From above, we know $c < c^*$ implies $IC_i - M$ is satisfied. Thus, we have a PE. Since the threshold is unique, there is no other full effort PE. Finally, since $c > c^*$ implies that $IC_i - M$ will fail to hold, so there is no full effort PE when $c > c^*$.

Finally, for completeness, note that at $c = c^*$the $IC_i - M$ condition holds with equality and $IC_p - M$ is satisfied with strict inequality, so we have a PE. □

Proof of Lemma 15. We simply verify that a unique $\bar{u}$ satisfies $IC - L$:

$$\bar{u}[1 + \delta p - \delta F(\bar{u})] - \delta \int_{\bar{u}}^{\infty} uf(u)du = 0.$$

The left-hand side is strictly increasing in $\bar{u}$. At $\bar{u} = 0$, the left-hand side is negative and equals $-\delta\mu$. As $\bar{u} \to \infty$ the second term dominates and the left-hand side goes to ∞. Thus, the left-hand side will cross zero, from below, exactly once. Let $\bar{u}^*_{SW}$ denote this unique $\bar{u}$ value. To verify the upper bound on $\bar{u}^*_{SW}$, we proceed as follows. First, note that the left-hand side above is bounded below:

$$\bar{u}[1 + \delta p - \delta F(\bar{u})] - \delta \int_{\bar{u}}^{\infty} uf(u)du > -\delta\mu + \bar{u}[1 - \delta(1 - p)],$$

where the linear lower bound is constructed by extrapolating from the left-hand side value of $-\delta\mu$ at $\bar{u} = 0$ at slope $1 - \delta(1 - p)$, which is below the slope of the left-hand side. If we show that this lower bound is above 0 at $\bar{u} = \mu/p$, then we know the left-hand side must cross 0 before $\bar{u} = \mu/p$ and, hence, that $\bar{u}^*_{SW} < \mu/p$. Calculating, we have

$$-\delta\mu + (\mu/p)[1 - \delta(1 - p)] = \frac{\mu(1 - \delta)}{p} > 0$$

and we are done. $\square$

Proof of Proposition 23. From Lemma 15 that $IC-L$ is satisfied if and only if the threshold is $\bar{u}^*_{SW}$. From the text, we know that $IC-M$ reduces to two cases. If $p\mu-c>0$ then IC_i-M is sufficient for IC_p-M to hold. We claim that IC_i-M holds if and only if $c^*<c<p\mu$, for some c^*. IC_i-M requires

$$p\bar{\mu}-c\geq\left[\frac{1+\delta p-\delta F(\bar{u})}{\delta p[1-F(\bar{u})]}\right](p\mu-c) \qquad \Leftrightarrow$$

$$\frac{c}{p}+\left[\frac{\delta p[1-F(\bar{u})]}{1+\delta p-\delta F(\bar{u})}\right]\left[E(\bar{u})-\frac{c}{p}\right]\geq\mu.$$

As we know from the proof of proposition 22, the left-hand side above is strictly increasing in c, below μ at $c=0$, and above μ at $c=p\mu$. Thus, there exists a unique c^* where the value of the left-hand side crosses μ. For the PE with initial shirking, this c^* corresponds to the value when $\bar{u}=\bar{u}^*_{SW}$. Given the case of $p\mu-c>0$, we have IC_i-M satisfied if $c^*\leq c\leq p\mu$ and we have a unique PE with initial shirking, and the cutoff is $\bar{u}^*_{SW}$. If $c<c^*$ then IC_i-M fails and no PE with initial shirking exists.

The second case is $p\mu-c<0$. Now, IC_p-M is sufficient for IC_i-M to hold. We claim that IC_p-M holds if and only if $p\mu<c<c^{**}$, for some c^{**}. IC_p-M requires that

$$p\bar{\mu}-c=pE(\bar{u})-c\geq 0$$

hold at the threshold $\bar{u}=\bar{u}^*_{SW}$, where we recall the definition $\bar{\mu}=E(\bar{u})=\int_{\bar{u}}^{\infty}uf(u)du/[1-$

$F(\bar{u})]$. Recall also that $E(\bar{u})$ is strictly increasing, that at $\bar{u} = 0$ we have $E(0) = \mu$, and that $E(\bar{u}) > \bar{u}$ holds for all $\bar{u}$. So, at $c = p\mu$, we have $pE(\bar{u}^*_{SW}) - c > pE(0) - c = 0$. As c rises, there is a unique $c^{**} > p\mu$, namely $c^{**} = pE(\bar{u}^*_{SW})$, where $IC_p - M$ will become binding. Thus, $IC_p - M$ holds for c between $p\mu$ and c^{**}. In this case, we have a unique PE with initial shirking, and the cutoff is $\bar{u}^*_{SW}$. But $IC_p - M$ fails for c above c^{**} and the equilibrium with initial shirking does not exist. □

Proof of Corollary 5. Recall that $\bar{u}^*_{SW}$ is the unique solution to

$$\delta \int_{\bar{u}}^{\infty} uf(u)du - \bar{u}[1 + \delta p - \delta F(\bar{u})] = 0,$$

where the left-hand side is strictly decreasing from $\delta\mu$ at $\bar{u} = 0$ to a limit of $-\infty$ as $\bar{u} \to \infty$. It is also convex since $\delta f(\bar{u}) > 0$. As a result, the left-hand side is bounded below by $\delta\mu - \bar{u}(1 + \delta p)$, where we are employing the slope at $\bar{u} = 0$, and it is bounded above by $\delta\mu - \bar{u}(1 + \delta p - \delta)$, where we are employing the limiting slope as $\bar{u} \to \infty$. Since the zero-crossing of the lower bound must be smaller than $\bar{u}^*_{SW}$, we have

$$p\mu < \frac{\delta\mu}{1 + \delta p} < \bar{u}^*_{SW} \qquad \Leftrightarrow \qquad \frac{p}{1 - p^2} < \delta$$

as the sufficient condition for $p\mu < \bar{u}^*_{SW}$. Note that $p/(1 - p^2)$ is increasing and convex, rising from 0 at $p = 0$ to ∞ as p goes to 1. By similar logic, the upper bound crosses zero to the left of $\bar{u}^*_{SW}$ and we have

$$\bar{u}^*_{SW} < \frac{\delta\mu}{1+\delta p-\delta} < p\mu \qquad \Leftrightarrow \qquad \frac{p}{1+p-p^2} > \delta$$

as the sufficient condition for $p\mu > \bar{u}^*_{SW}$. Note that $p/(1+p-p^2)$ is increasing, initially concave and then convex, rising from 0 at $p=0$ to 1 as p goes to 1. Since p is between 0 and 1, we always have $p/(1-p^2) > p/(1+p-p^2)$.

Recall that part (iii) was demonstrated in the proof of Lemma 15. For part (iv), we must show $\bar{u}^*_{SW} < \bar{u}^*_{WW}$. Recall that $\bar{u}^*_{WW}$ is the unique solution to

$$\bar{u}[1-\delta(1-p)F(\bar{u})]-\delta(1-p)\int_{\bar{u}}^{\infty} uf(u)du = [1-\delta(1-p)]\mu, \tag{61}$$

and that the left-hand side is concave and strictly increasing from $-\delta(1-p)\mu$ at $\bar{u}=0$ to ∞ as $\bar{u}\to\infty$. If we can show that the left-hand side is below $[1-\delta(1-p)]\mu$ at $\bar{u}=\bar{u}^*_{SW}$, then we must have $\bar{u}^*_{SW} < \bar{u}^*_{WW}$ since the left-hand side must cross $[1-\delta(1-p)]\mu$ to the right of $\bar{u}^*_{SW}$. We know that $\bar{u}^*_{SW}$ satisfies

$$\bar{u}^*_{SW}[1+\delta p-\delta F(\bar{u}^*_{SW})] = \delta\int_{\bar{u}^*_{SW}}^{\infty} uf(u)du$$

and, upon substitution in the left-hand side of (61), we have

$$\begin{aligned}
\bar{u}^*_{SW}[1-\delta(1-p)F(\bar{u}^*_{SW})]-\delta(1-p)\int_{\bar{u}^*_{SW}}^{\infty} uf(u)du \lesseqgtr [1-\delta(1-p)]\mu &\quad \Leftrightarrow \\
\bar{u}^*_{SW}\left\{[1-\delta(1-p)F(\bar{u}^*_{SW})]-(1-p)[1+\delta p-\delta F(\bar{u}^*_{SW})]\right\} \lesseqgtr [1-\delta(1-p)]\mu &\quad \Leftrightarrow \\
\bar{u}^*_{SW}p[1-\delta(1-p)] \lesseqgtr [1-\delta(1-p)]\mu &\quad \Leftrightarrow \\
\bar{u}^*_{SW} \lesseqgtr \mu/p. &
\end{aligned}$$

By part (iii), we have $\bar{u}^*_{SW} < \mu/p$ and, hence, the left-hand side is below $[1-\delta(1-p)]\mu$ at $\bar{u} = \bar{u}^*_{SW}$ and we are done. □

Proof of Proposition 24. Given the discussion in the text, we need only show that

$IC_p - M$, $0 \geq p\bar{\mu} - c$, holds for $\bar{u} \leq \bar{u}^*_0$ and fails for larger $\bar{u}$. Recall once again our definition for the conditinal mean $\bar{\mu} = E(\bar{u}) = \int_{\bar{u}}^{\infty} uf(u)du/[1-F(\bar{u})]$ and that $E(\bar{u})$ is strictly increasing, with $E(0) = \mu$ at $\bar{u} = 0$, and that $E(\bar{u}) > \bar{u}$ holds for all $\bar{u}$. The bound $\bar{u}^*_0$ is the unique solution to

$$\frac{\int_{\bar{u}}^{\infty} uf(u)du}{1-F(\bar{u})} = \frac{c}{p}.$$

By monotonicity, we have $E(\bar{u}) < c/p$ for $\bar{u} < \bar{u}^*_0$ and, consequently, $IC_p - M$ holds for all $\bar{u} \leq \bar{u}^*_0$. On the other side, again by monotonicy, when $\bar{u} > \bar{u}^*_0$ we see that $IC_p - M$ fails. Finally, since $E(\bar{u}) > \bar{u}$ holds at any $\bar{u}$, we have $c/p = E(\bar{u}^*_0) > \bar{u}^*_0$. □

Proof of Proposition 25. There are two parts to the proof. First, we must prove existence by constructing a full-effort PE in which $\bar{u}$ is identical to the one-manager case. Second, we must prove uniqueness by showing that this is the only possible full-effort PE.

Begin by supposing that each manager works iff $u > c/p$ and that L persists iff $u > \bar{u}$, where $\bar{u} \in (c/p, \infty)$. This is a PE iff two conditions hold:

1. (Leader IC) $p\bar{u} = (1-\delta)V_L$.
2. (Manager IC) $\bar{u} \geq c/p - \delta\left[V_M - \frac{c}{2(1-\delta)}\right]$.

The leader's value V_L in such a PE is given by

$$V_L = \int_0^{c/p} \delta V_L f(u)du + \int_{c/p}^{\bar{u}} (pu + \delta V_L) f(u)du + \int_{\bar{u}}^{\infty} \frac{p(u+\delta V_L)}{1-\delta(1-p)} f(u)du.$$

Solving for V_L yields

$$V_L = \frac{\int_{c/p}^{\infty} puf(u)du - \delta(1-p)\int_{c/p}^{\bar{u}} puf(u)du}{(1-\delta)(1-\delta(1-p)F(\bar{u}))}.$$

From this solution, we see that Leader IC holds iff

$$\bar{u}[1-\delta(1-p)F(\bar{u})] + \delta(1-p)\int_{c/p}^{\bar{u}} uf(u)du = \int_{c/p}^{\infty} uf(u)du. \tag{62}$$

Direct inspection reveals that (62) admits a unique solution that is above c/p whenever (59) holds.

To finish the proof, we must show the Manager IC is satisfied whenever (59) holds. To do this, first solve for V_M in a similar way to how we solved for V_L. After simplification, the solution for V_M can be written as

$$V_M = \frac{\int_{c/p}^{\infty}(pu - c/2)f(u)du + \delta(1-p)\int_{c/p}^{\bar{u}}(pu - c/2)f(u)du}{(1-\delta)(1-\delta(1-p)F(\bar{u}))}. \tag{63}$$

Observe that V_M is equal to V_L minus the expected PDV costs of effort. Hence, the Manager IC is equivalent to

$$\bar{u} \geq \frac{c}{p} - \delta\left[V_L - \frac{c}{2(1-\delta)}\left(2 - \frac{(1-\delta(1-p))F(c/p)}{1-\delta(1-p)F(\bar{u})}\right)\right].$$

Since $p\bar{u} = (1-\delta)V_L$ in any high-$\bar{u}$ equilibrium, the bracketed term on the right-hand side of Manager IC is positive whenever $\bar{u} > c/p$. From this, we conclude that Manager IC is always satisfied whenever (59) holds. □

Proof of Proposition 26. We begin by showing that a full-effort PE does not exist. To do this, assume by way of contradiction that a full-effort PE does exist. Then we know that the leader's choice of $\bar{u}$ must satisfy $\bar{u} < c/p$ because (1) fails. Hence, each manager's value in this PE satisfies

$$V_M = \int_0^{\bar{u}} \delta V_M f(u)du + \int_{\bar{u}}^{\infty} \frac{p(u+\delta V_L) - c/2}{1-\delta(1-p)} f(u)du,$$

which solves to give

$$V_M = \frac{\int_{\bar{u}}^{\infty}(pu - c/2)f(u)du}{(1-\delta)(1-\delta(1-p)F(\bar{u}))}.$$

As before, the value V_M is equal to V_L minus the expected PDV costs of effort. Using this identity, we see that Manager IC is equivalent to

$$\bar{u} \geq \frac{c}{p} - \delta \left[V_L - \frac{c}{2(1-\delta)} \left(1 + \frac{1-F(\bar{u})}{1-\delta(1-p)F(\bar{u})} \right) \right]. \tag{64}$$

Since $\frac{1-F(\bar{u})}{1-\delta(1-p)F(\bar{u})} > 1$ and $V_L = \frac{p\bar{u}}{1-\delta} < \frac{c}{1-\delta}$, the bracketed term on the right-hand side of (64) is negative whenever $\bar{u} < c/p$. From this, we conclude that (64) fails whenever the future is dim. Hence, a full-effort PE cannot exist, and shirking must occur over a non-trivial subset of $(\bar{u}, \infty)$.

To characterize the set of projects over which shirking occurs, recall that manager i's IC condition given $\sigma_L = 1$ and $e_j = 1$ is

$$u \geq \frac{c}{p} - \delta \left[V_i - \frac{c}{2(1-\delta)} \right].$$

We claim that the right-hand side term is greater than c/p. To see this, first note that $V_i < V_L < \frac{p\bar{u}}{2(1-\delta)}$ holds in any shirking PE. Then, combine this observation with the fact that $\bar{u} < c/p$ whenever the future is dim to conclude that $V_i - \frac{c}{2(1-\delta)}$ must be negative. As a result, the Manager IC given $\sigma_L = 1$ and $e_j = 1$ holds only if $u \geq u^*$, where $u^* \in (c/p, \infty)$ solves (64) as an equality.

To complete the proof, we must argue that L has no incentive to terminate projects with value $u > c/p$ when the future is dim, even if one manager shirks. To do this, note that the leader's value in the proposed PE satisfies

$$V_L = \int_0^{\bar{u}} \delta V_L f(u)du + \int_{\bar{u}}^{u^*} \frac{(p/2)(u+\delta V_L)}{1-\delta(1-p/2)} f(u)du + \int_{u^*}^{\bar{u}} \frac{p(u+\delta V_L)}{1-\delta(1-p)} f(u)du,$$

Solving this equation for V_L yields the solution

$$V_L = \frac{\int_{\bar{u}}^{\infty} puf(u)du - (1-\delta(1-2p)) \int_{\bar{u}}^{u_2^*} (p/2)uf(u)du}{(1-\delta(1-p/2))(1-\delta(1-p))G(\bar{u},u_2^*)},$$

where

$$G(x,y) = 1-\delta F(x) - \frac{(p/2)\delta}{1-\delta(1-p/2)}(F(y)-F(x)) - \frac{p\delta}{1-\delta(1-p)}(1-F(u_2^*)).$$

Using this solution, we can simplify the leader's IC condition to obtain

$$\bar{u}G(\bar{u},u_2^*) - \frac{1}{1-\delta(1-p/2)} \int_{\bar{u}}^{u_2^*} uf(u)du \geq \frac{1-\delta}{1-\delta(1-p)} \int_{u_2^*}^{\infty} uf(u)du \qquad (65)$$

The right-hand side (65) is constant in $\bar{u}$, while the derivative of the left-hand side is

$$\frac{\partial}{\partial \bar{u}} \left\{ \bar{u}G(\bar{u},u_2^*) - \frac{1}{1-\delta(1-p/2)} \int_{\bar{u}}^{u_2^*} uf(u)du \right\}$$

$$= G(\bar{u},u_2^*) + \left[1 - \frac{(1-p/2)(1-\delta)\delta}{1-\delta(1-p/2)} \right] \bar{u}f(u)$$

This derivative is positive; hence, there is at most one solution to (65) an equality.

We wish to strengthen this conclusion by showing that a solution exists. We do this by inspecting limits as $\bar{u} \to 0$ and as $\bar{u} \to u_2^*$. As $\bar{u} \to 0$, the left-hand side of (2) approaches zero, while the left-hand side remains positive. Thus, the leader's IC fails whenever $\bar{u}$ is close to zero.

As $\bar{u} \to \bar{u}_2^*$, the left-hand side of (2) approaches $u_2^*\left[1-\delta F\left(u_2^*\right)-\frac{p\delta}{1-\delta(1-p)}(1-F(u_2^*))\right]$, while the right-hand side approaches $\frac{1-\delta}{1-\delta(1-p)}\int_{u_2^*}^{\infty} uf(u)du$. Thus, a necessary condition and sufficient condition for Leader IC to hold at some $\bar{u} \in (0, \bar{u}_2^*)$ is that

$$\bar{u}_2^*\left[(1-\delta(1-p))(1-\delta F(u_2^*))-p\delta(1-F(u_2^*))\right] > (1-\delta)\int_{u_2^*}^{\infty} uf(u)du.$$

Upon simplification, this inequality reduces to

$$u_2^*\left[1-\delta(1-p)F(u_2^*)\right] > \int_{u_2^*}^{\infty} uf(u)du.$$

The above inequality is analogous to the dim future condition from the full-effort PE case. It shows that a solution to Leader IC holds for some $\bar{u} \in (0, \bar{u}_2^*)$; in particular, $\bar{u}$ may be above or below c/p. If $\bar{u} < c/p$, then both managers work for all $u \in (u^*, \infty)$ and only one manager works if $u \in (\bar{u}, u^*)$. If $\bar{u} > c/p$, then a low-$\bar{u}$ shirking PE is impossible. Thus, the only possibility is $\bar{u} > c/p$. In this case, both managers work for all $u \in (c/p, \bar{u}) \cup (u^*, \infty)$ and only one manager works if $\bar{u} \in (\bar{u}, u^*)$. □

Chapter 4

Conclusion

In this dissertation, we developed and analyzed three models of innovation. We have shown the central implications of three sources of endogenous incentives on equilibrium investment behavior. We showed in Chapter 2 that endogenous learning plays a critical role in the timing of innovation adoption. Furthermore, we explored the interaction between endogenous learning and market asymmetric to highlight stylized empirical relationships between industry dominance, namely ex-ante quality, and subsequent adoption decisions. We showed in Chapter 3 that endogenous product market incentive generate substantial effects in the classic patent race. Finally, we showed in Chapter 3 that endogenous persistence plays a critical role in the process of innovation within an organization, and can be used as a valuable tool to overcome agency conflicts between leaders and managers. Each of the models highlights the importance of endogenous mechanisms at key stages of the innovation process. By highlighting these endogenous mechanisms, this dissertation aims to improve our understanding of the innovation process so that we may better design innovation policy, better management of research and development into innovation, and, ultimately, maximize economic growth and prosperity.

References

Acemoglu, D., K. Bimpikis, and A. Ozdaglar (2011): “Experimentation, Patents, and Innovation,” *American Economic Journal: Microeconomics*, 3, 37–77.

Aghion, P., P. Bolton, C. Harris, and B. Jullien (1991): “Optimal Learning by Experimentation,” *Review of Economic Studies*, 58, 621–654.

Aghion, P. and J. Tirole (1997): “"Formal and Real Authority in Organizations",” *Journal of Political Economy*, 105, 1–29.

Akcigit, U. and Q. Liu (2016): “The Role of Information in Innovation and Competition,” *Journal of the European Economic Association*, 14, 828–870.

Andries, P. and P. Hünermund (2020): “Firm-level effects of staged investments in innovation: The moderating role of resource availability,” *Research Policy*, 49, 1–14.

Anton, J. and D. Yao (2022): “Delay as Agenda Setting,” *Working Paper*.

Anton, M., F. Ederer, M. Gine, and M. C. Schmalz (2021): “Common Ownership, Competition, and Top Management Incentives,” *Working Paper*.

Arrow, K. (1962): “Economic Welfare and the Allocation of Resources for Invention,” in *The Rate and Direction of Inventive Activity: Economic and Social Factors*, Princeton University Press, 609–626.

Athey, S. and A. Schmutzler (2001): “Investment and Market Dominance,” *The RAND Journal of Economics*, 32, 1–26.

Awaya, Y. and V. Krishna (2019): “Startups and Upstarts,” SSRN Scholarly Paper ID 3380351, Social Science Research Network, Rochester, NY.

——— (2021): “Startups and Upstarts: Disadvantageous Information in R&D,” *Journal of Political Economy*, 129, 534–569.

Bagwell, K. and G. Ramey (1994): “Coordination Economies, Advertising, and Search Behavior in Retail Markets,” *American Economic Review*, 84, 498–517.

Baker, J. B. (1995): “Fringe Firms and the Incentive to Innovate,” *Antitrust Law Journal*, 63, 621–641.

Baker, George, R. G. and K. Murphy (2002): “Relational Contracts and the Theory of the Firm,” *Quarterly Journal of Economics*, 117, 39–84.

Balachandra, R. and J. Raelin (1980): “How to Decide When to Abandon a Project,” *Research Management*, 23, 24–29.

Becker, G. S. (1965): "A Theory of the Allocation of Time," *The Economic Journal*, 75, 493–517.

Bergemann, D. and U. Hege (2005): "The Financing of Innovation: Learning and Stopping," *RAND Journal of Economics*, 36, 719–752.

Bergemann, D. and J. Välimäki (1997): "Market Diffusion with Two-Sided Learning," *RAND Journal of Economics*, 28, 773–795.

——— (2917): "Bandit Problems," in *The New Palgrave Dictionary of Economics*, 336–340.

Bergemann, D. and J. Välimäki (2000): "Experimentation in Markets," *The Review of Economic Studies*, 67, 213–234.

——— (2002): "Entry and Vertical Differentiation," *Journal of Economic Theory*, 106, 91–125.

Bimpikis, K., S. Ehsani, and M. Mostagir (2019): "Designing Dynamic Contests," *Operations Research*, 67, 339–356.

Blundell, R., R. Griffith, and J. Van Reenen (1999): "Market Share, Market Value and Innovation in a Panel of British Manufacturing Firms," *The Review of Economic Studies*, 66, 529–554.

Bolton, P. and C. Harris (1999): "Strategic Experimentation," *Econometrica*, 67, 349–374.

——— (2000): "Strategic experimentation: The undiscounted case," *Incentives, Organizations and Public Economics–Papers in Honour of Sir James Mirrlees*, 53–68.

Bonatti, A. (2011): "Menu Pricing and Learning," *American Economic Journal: Microeconomics*, 3, 124–163.

Bourreau, M. and B. Jullien (2018): "Mergers, investments and demand expansion," *Economics Letters*, 167, 136–141.

Bower, J. (1970): *Managing the Resource Allocation Process: A Study of Corporate Planning and Investment*, Harvard Business School, Division of Research.

Bower, J. and C. Gilbert (2005): "A Revised Model of the Resource Allocation Process," in *From Resource Allocation to Strategy*, Oxford University Press.

Brennan, Tom, P. E. J. K. and E. Roth (2020): "Building an R&D Strategy for Modern Times," *McKinsey Insights*.

Brown, J., G. Martinsson, and B. Petersen (2012): "Do financing constraints matter for R&D?" *European Economic Review*, 56, 1512–1529.

Brunner, David, L. F. A. M. and D. Zinner (2008): “R&D Project Selection and Portfolio Management: A Review of the Past, a Description of the Present, and a Sketch of the Future,” in *Handbook of Technology and Innovation Management*, Wiley, 215–238.

Bryan, K. A. and J. Lemus (2017): “The direction of innovation,” *Journal of Economic Theory*, 172, 247–272.

Buell, C. (1967): “When to Terminate a Research and Development Project,” *Research Management*, 10, 275–284.

Burgelman, R. (1991): “Intraorganizational Ecology of Strategy Making and Organizational Adaptation: Theory and Field Research,” *Organization Science*, 2, 239–262.

Cabral, L. M. B. (2018): “Standing on the Shoulders of Dwarfs: Dominant Firms and Innovation Incentives,” Tech. Rep. 13115, C.E.P.R. Discussion Papers.

Cabral, L. M. B. and M. H. Riordan (1994): “The Learning Curve, Market Dominance, and Predatory Pricing,” *Econometrica*, 62, 1115–1140.

Chan, Tat, J. N. and H. Owan (2007): “Strategic Management of R&D Pipelines with Cospecialized Investments and Technology Markets,” *Management Science*, 53, 667–682.

Chandy, R. K. and G. J. Tellis (2000): “The Incumbent’s Curse? Incumbency, Size, and Radical Product Innovation,” *Journal of Marketing*, 64, 1–17.

Choi, J. P. (1991): “Dynamic R&D Competition under "Hazard Rate" Uncertainty,” *The RAND Journal of Economics*, 22, 596–610.

Christensen, C. (1997): *The Innovator’s Dilemma: When New Technologies Cause Great Firms to Fail*, Boston, MA: Harvard Business School Press.

Christensen, C. M. (1993): “The Rigid Disk Drive Industry: A History of Commercial and Technological Turbulence,” *Business History Review*, 67, 531–588.

Cohen, W. M. (2010a): “Chapter 4 - Fifty Years of Empirical Studies of Innovative Activity and Performance,” in *Handbook of the Economics of Innovation*, ed. by B. H. Hall and N. Rosenberg, North-Holland, vol. 1 of *Handbook of The Economics of Innovation, Vol. 1*, 129–213.

——— (2010b): “Fifty Years of Empirical Studies of Innovative Activity and Performance,” in *Handbook of the Economics of Innovation*, Elsevier, vol. 1, 129–213.

Cohen, W. M. and S. Klepper (1996): “Firm Size and the Nature of Innovation within Industries: The Case of Process and Product R&D,” *The Review of Economics and Statistics*, 78, 232–243.

Cohen, W. M. and R. C. Levin (1989): “Chapter 18 Empirical studies of innovation and market structure,” in *Handbook of Industrial Organization*, Elsevier, vol. 2, 1059–1107.

Comin, D. and B. Hobijn (2010): “An Exploration of Technology Diffusion,” *American Economic Review*, 100, 2031–2059.

Cooper, A. C. and D. Schendel (1976): “Strategic responses to technological threats,” *Business Horizons*, 19, 61–69.

Cooper, Robert, S. E. and E. Kleinschmidt (1998): “Best Practices for Managing R&D Portfolios,” *Research Technology Management*, 41, 20–33.

Cripps, M. and C. Thomas (2019): “Strategic experimentation in queues,” *Theoretical Economics*, 14, 647–708.

Cunningham, C., F. Ederer, and S. Ma (2021): “Killer Acquisitions,” *Journal of Political Economy*, 129, 649–702.

Damanpour, F. and S. Gopalakrishnan (2001): “The Dynamics of the Adoption of Product and Process Innovations in Organizations,” *Journal of Management Studies*, 38, 45–65.

Das, K. and N. Klein (2020a): “Do Stronger Patents Lead to Faster Innovation? The Effect of Duplicative Search,” *Working Paper*, 28.

——— (2020b): “Over-and Under-Experimentation in a Patent Race with Private Learning,” *Working Paper*, 24.

Décamps, J.-P. and T. Mariotti (2004): “Investment timing and learning externalities,” *Journal of Economic Theory*, 118, 80–102.

Denicolo, V. and M. Polo (2018): “The Innovation Theory of Harm: An Appraisal,” *Antitrust Law Journal*, 82, 921.

Dewar, R. D. and J. E. Dutton (1986): “The Adoption of Radical and Incremental Innovations: An Empirical Analysis,” *Management Science*, 32, 1422–1433.

Doraszelski, U. (2003): “An R&D Race with Knowledge Accumulation,” *The RAND Journal of Economics*, 34, 20–42.

Dosi, G. (1988): “Sources, Procedures, and Microeconomic Effects of Innovation,” *Journal of Economic Literature*, 26, 1120–1171.

Dutta, P. K. (1991): “What do discounted optima converge to?: A theory of discount rate asymptotics in economic models,” *Journal of Economic Theory*, 55, 64–94.

Dutta, P. K., S. Lach, and A. Rustichini (1995): “Better Late than Early: Vertical Differentiation in the Adoption of a New Technology,” *Journal of Economics & Management*

Strategy, 4, 563–589.

Dutta, P. K. and A. Rustichini (1993a): "A theory of stopping time games with applications to product innovations and asset sales," *Economic Theory*, 3, 743–763.

——— (1993b): "A Theory of Stopping Time Games with Applications to Product Innovations and Asset Sales," *Economic Theory*, 3, 743–763.

Décamps, J.-P., F. Gensbittel, and T. Mariotti (2021): "Investment Timing and Technological Breakthroughs," .

Fagerberg, J., D. C. Mowery, and J. Cantwell (2009): *Innovation and Competitiveness*, Oxford University Press.

Farrell, J. and C. Shapiro (1990): "Horizontal Mergers: An Equilibrium Analysis," *The American Economic Review*, 80, 107–126.

Federico, G., F. S. Morton, and C. Shapiro (2019): "Antitrust and Innovation: Welcoming and Protecting Disruption," in *NBER Chapters*, National Bureau of Economic Research, Inc, 125–190.

——— (2020): "Antitrust and Innovation: Welcoming and Protecting Disruption," *Innovation Policy and the Economy*, 20, 125–190.

Frick, M. and Y. Ishii (2015): "Innovation Adoption by Forward-Looking Social Learners," SSRN Scholarly Paper ID 2565769, Social Science Research Network, Rochester, NY.

Fudenberg, D. and J. Tirole (1985): "Preemption and Rent Equalization in the Adoption of New Technology," *The Review of Economic Studies*, 52, 383–401.

Ghemawat, P. (1991): *Commitment: The Dynamic of Strategy*, New York: Free Press.

Giebel, M. and K. Kraft (2019): "External Financing Constraints and Firm Innovation," *The Journal of Industrial Economics*, 67, 91–126.

Gifford, S. (1992): "Allocation of entrepreneurial attention," *Journal of Economic Behavior & Organization*, 19, 265–284.

Gilbert, C. (2006a): "Change in the Presence of Residual Fit: Can Competing Frames Coexist?" *Organization Science*, 17, 150–167.

Gilbert, J. T. and P. H. Birnbaum-More (1996): "Innovation timing advantages: From economic theory to strategic application," *Journal of Engineering and Technology Management*, 12, 245–266.

Gilbert, R. (2006b): "Looking for Mr. Schumpeter: Where Are We in the Competition–Innovation Debate?" *Innovation Policy and the Economy*, 6, 159–215.

——— (2006c): "Looking for Mr. Schumpeter: Where Are We in the Competition-Innovation Debate?" in *Innovation Policy and the Economy, Volume 6*, National Bureau of Economic Research, Inc, 159–215.

Gilbert, R. and C. Shapiro (1990): "Optimal Patent Length and Breadth," *The RAND Journal of Economics*, 21, 106–112.

Golder, P. N. and G. J. Tellis (1993): "Pioneer Advantage: Marketing Logic or Marketing Legend?" *Journal of Marketing Research*, 30, 158–170.

Gowrisankaran, G. and J. Stavins (2004): "Network Externalities and Technology Adoption: Lessons from Electronic Payments," *The RAND Journal of Economics*, 35, 260–276.

Green, J. R. and S. Scotchmer (1995): "On the Division of Profit in Sequential Innovation," *The RAND Journal of Economics*, 26, 20–33.

Guler, I. (2007): "Throwing Good Money after Bad? Political and Institutional Influences on Sequential Decision Making in the Venture Capital Industry," *Administrative Science Quarterly*, 52, 248–285.

——— (2018): "Pulling the Plug: The Capability to Terminate Unsuccessful Projects and Firm Performance," *Strategy Science*, 3, 481–497.

Halac, M., N. Kartik, and Q. Liu (2016): "Optimal Contracts for Experimentation," *The Review of Economic Studies*, 83, 1040–1091.

Halac, Marina, N. K. and Q. Liu (2016): "Optimal Contracts for Experimentation," *Review of Economic Studies*, 83, 1040–1091.

Hall, B. H. and B. Khan (2003): "Adoption of New Technology," Tech. Rep. 9730, National Bureau of Economic Research, Inc.

Haneda, S. and A. Ono (2022): *R&D Management Practices and Innovation: Evidence from a Firm Survey*, Springer (Open Access).

Hannan, T. and J. McDowell (1984): "The Determinants of Technology Adoption: The Case of the Banking Firm," *RAND Journal of Economics*, 15, 328–335.

Heclo, H. (1977): *A Government of Strangers: Executive Politics in Washington*, Brookings.

Heidhues, P., S. Rady, and P. Strack (2015): "Strategic experimentation with private payoffs," *Journal of Economic Theory*, 159, 531–551.

Helfat, C. and K. Eisenhardt (2004): "Inter-temporal economies of scope, organizational modularity, and the dynamics of diversification," *Strategic Management Journal*, 25, 1217–1232.

Hellman, T. and V. Thiele (2011): "Incentives and Innovation: A Multitasking Approach," *American Economic Journal: Microeconomics*, 3, 78–128.

Henderson, R. and K. Clark (1990): "Architectural Innovation: The Reconfiguration of Existing Product Technologies and the Failure of Established Firms," *Administrative science quarterly*, 35.

Hendricks, K. (1992): "Reputations in the Adoption of a New Technology," *International Journal of Industrial Organization*, 663–677.

Hendricks, K., A. Weiss, and C. Wilson (1988): "The War of Attrition in Continuous Time with Complete Information," *International Economic Review*, 29, 663–680.

Hermalin, B. (1998): "Toward an Economic Theory of Leadership: Leading by Example," *American Economic Review*, 88, 1188–1206.

Holmes, T. J., D. K. Levine, and J. A. Schmitz (2012): "Monopoly and the Incentive to Innovate When Adoption Involves Switchover Disruptions," *American Economic Journal: Microeconomics*, 4, 1–33.

Hopenhayn, H., G. Llobet, and M. Mitchell (2006): "Rewarding Sequential Innovators: Prizes, Patents, and Buyouts," *Journal of Political Economy*, 114, 1041–1068.

Hopenhayn, H. and F. Squintani (2021): "On the Direction of Innovation," *Journal of Political Economy*, 129, 1991–2022.

Hoppe, H. C. (2000): "Second-Mover Advantages in the Strategic Adoption of New Technology under Uncertainty," *International Journal of Industrial Organization*, 18, 315–338.

——— (2002): "The Timing of New Technology Adoption: Theoretical Models and Empirical Evidence," *The Manchester School*, 70, 56–76.

Horner, J. and A. Skrzypacz (2017): "Learning, Experimentation, and Information Design," *Advances in Economics and Econometrics*, 1, 63–98.

Hörner, J. and A. Skrzypacz (2017): "Learning, Experimentation, and Information Design," in *Advances in Economics and Econometrics: Eleventh World Congress*, ed. by A. Pakes, B. Honoré, L. Samuelson, and M. Piazzesi, Cambridge: Cambridge University Press, vol. 1 of *Econometric Society Monographs*, 63–98.

Igami, M. (2017): "Estimating the Innovator's Dilemma: Structural Analysis of Creative Destruction in the Hard Disk Drive Industry, 1981-1998," *Journal of Political Economy*, 125, 798–847.

Jekunen, A. (2014): “Decision-making in Product Portfolios of Pharmaceutical Research and Development: Managing Streams of Innovation in Highly Regulated Markets,” *Drug Design, Development, and Therapy*, 8, 2009–2016.

Jensen, R. (1982): “Adoption and diffusion of an innovation of uncertain profitability,” *Journal of Economic Theory*, 27, 182–193.

——— (1992): “Innovation Adoption and Welfare Under Uncertainty,” *The Journal of Industrial Economics*, 40, 173–180.

Jensen, R. A. (1980): “A Duopoly Model of the Adoption of an Innovation of Uncertainty Profitability,” Tech. Rep. 434, Northwestern University, Center for Mathematical Studies in Economics and Management Science.

Judd, K. L. and M. H. Riordan (1994): “Price and Quality in a New Product Monopoly,” *Review of Economic Studies*, 61, 773–789.

Jullien, B. and Y. Lefouili (2018): “Horizontal Mergers and Innovation,” *Journal of Competition Law & Economics*, 14, 364–392.

Kapur, S. (1995): “Technological Diffusion with Social Learning,” *The Journal of Industrial Economics*, 43, 173–195.

Katz, M. and C. Shapiro (1987): “R&D Rivalry with Licensing or Imitation,” *American Economic Review*, 77, 402–20.

Katz, M. L. and C. Shapiro (1985): “On the Licensing of Innovations,” *The RAND Journal of Economics*, 16, 504–520.

Keller, G. and S. Rady (2010): “Strategic experimentation with Poisson bandits,” *Theoretical Economics*, 5, 275–311.

——— (2015): “Breakdowns,” *Theoretical Economics*, 10, 175–202.

——— (2020): “Undiscounted bandit games,” *Games and Economic Behavior*, 124, 43–61.

Keller, G., S. Rady, and M. Cripps (2005): “Strategic Experimentation with Exponential Bandits,” *Econometrica*, 73, 39–68.

Khanna, Rajat, I. G. and A. Nerkar (2018): “Entangled decisions: Knowledge interdependencies and terminations of patented inventions in the pharmaceutical industry,” *Strategic Management Journal*, 39, 2439–2465.

Klein, N. and S. Rady (2011): “Negatively Correlated Bandits,” *The Review of Economic Studies*, 78, 693–732.

Klepper, S. (1996): "Entry, Exit, Growth, and Innovation over the Product Life Cycle," *American Economic Review*, 86, 562–583.

Klingebiel, R. and J. Joseph (2016): "Entry timing and innovation strategy in feature phones," *Strategic Management Journal*, 37, 1002–1020.

Kremer, M. (1998): "Patent Buyouts: A Mechanism for Encouraging Innovation*," *The Quarterly Journal of Economics*, 113, 1137–1167.

Krieger, J. (2021): "Trials and Terminations: Learning from Competitors' R&D Failures," *Management Science*, 67, 5525–5548.

Krylov, N. V. (1980): *Controlled Diffusion Processes*, Stochastic Modelling and Applied Probability, Berlin Heidelberg: Springer-Verlag.

Langer, A. and D. Lemoine (2022): "Designing Dynamic Subsidies to Spur Adoption of New Technologies," *Journal of the Association of Environmental and Resource Economists*, 9, 1197–1234.

Lavie, D., U. Stettner, and M. L. Tushman (2010): "Exploration and Exploitation Within and Across Organizations," *Academy of Management Annals*, 4, 109–155.

Levinthal, D. and B. Wu (2010): "Opportunity costs and non-scale free capabilities: profit maximization, corporate scope, and profit margins," *Strategic Management Journal*, 31, 780–801.

Lieberman, M. B. and D. B. Montgomery (1988): "First-Mover Advantages," *Strategic Management Journal*, 9, 41–58.

Liu, Q. and Y. F. Wong (2022): "Strategic Exploration: Preemption and Prioritization," *Working Paper*.

Majure, R. (2022): "Evaluating Innovation Theories of Harm in Merger Review: Economic Frameworks and Difficulties," .

Malueg, D. A. and S. O. Tsutsui (1997): "Dynamic R&D Competition with Learning," *The RAND Journal of Economics*, 28, 751–772.

Mamer, J. W. and K. F. McCardle (1987): "Uncertainty, Competition, and the Adoption of New Technology," *Management Science*, 33, 161–177.

Manso, G. (2011): "Motivating Innovation," *Journal of Finance*, 66, 1823–1860.

March, J. (1991): "Exploration and Exploitation in Organizational Learning," *Organization Science*, 2, 71–87.

Maskin, E. and J. Tirole (2001): "Markov Perfect Equilibrium: I. Observable Actions," *Journal of Economic Theory*, 100, 191–219.

McAfee, R. P. and M. A. Williams (1992): "Horizontal Mergers and Antitrust Policy," *The Journal of Industrial Economics*, 40, 181–187.

Miller, A. R. and C. Tucker (2014): "Health information exchange, system size and information silos," *Journal of Health Economics*, 33, 28–42.

Moscarini, G. and F. Squintani (2010): "Competitive experimentation with private information: The survivor's curse," *Journal of Economic Theory*, 145, 639–660.

Motta, M. (2004a): *Competition Policy: Theory and Practice*, Cambridge University Press.

——— (2004b): *Competition Policy. Theory and Practice*, Cambridge University Press.

Nocke, V. and M. D. Whinston (2010): "Dynamic Merger Review," *Journal of Political Economy*, 118, 1201–1251.

Office, C. B. (2021): *Research and Development in the Pharmaceutical Industry*, GPO.

Peck, Richard, D. L. I. G. C. L. and J. Isaacs (????): "Why is it hard to terminate failing projects in pharmaceutical R&D?" *Nature Reviews Drug Discovery*, 14, 663–664.

Peitz, M., S. Rady, and P. Trepper (2017): "Experimentation in Two-Sided Markets," *Journal of the European Economic Association*, 15, 128–172.

Radner, R. and M. Rothschild (1975): "On the allocation of effort," *Journal of Economic Theory*, 10, 358–376.

Rantakari, H. (2012): "Employee Initiative and Managerial Control," *American Economic Journal: Microeconomics*, 4, 171–211.

Reinganum, J. F. (1981): "On the Diffusion of New Technology: A Game Theoretic Approach," *Review of Economic Studies*, 48, 395–405.

——— (1982): "A Dynamic Game of R and D: Patent Protection and Competitive Behavior," *Econometrica*, 50, 671–688.

——— (1983): "Technology Adoption under Imperfect Information," *Bell Journal of Economics*, 14, 57–69.

——— (1989): "Chapter 14 The timing of innovation: Research, development, and diffusion," in *Handbook of Industrial Organization*, Elsevier, vol. 1, 849–908.

Riordan, M. H. and D. J. Salant (1994): "Preemptive Adoptions of an Emerging Technology," *Journal of Industrial Economics*, 42, 247–261.

Rivkin, Jan, D. L. and G. Hamel (2006): *Change at Whirlpool Corporation (A), (B), and*

(C), Harvard Business School Case 9-705-462, 463, 464.

Rosen, R. J. (1991): "Research and Development with Asymmetric Firm Sizes," *RAND Journal of Economics*, 22, 411–429.

Rosenbloom, R. and E. Pruyne (1997): "Polaroid Corporation: Digital Imaging Technology in 1997," *Harvard Business School Case*, 798-013.

Schivardi, F. and M. Schneider (2008): "Strategic experimentation and disruptive technological change," *Review of Economic Dynamics*, 11, 386–412.

Schlee, E. E. (1996): "The Value of Information about Product Quality," *The RAND Journal of Economics*, 27, 803–815.

Schumpeter, J. (1934): *The Theory of Economic Development*, London: Oxford University Press.

Shapiro, C. (2011): "Competition and Innovation: Did Arrow Hit the Bull's Eye?" in *The Rate and Direction of Inventive Activity Revisited*, University of Chicago Press, vol. 1, 361–404.

——— (2012): "Competition and Innovation: Did Arrow Hit the Bull's Eye?" in *The Rate and Direction of Inventive Activity Revisited*, ed. by J. Lerner and S. Stern, University of Chicago Press, 361–404.

Simon, L. K. and M. B. Stinchcombe (1989): "Extensive Form Games in Continuous Time: Pure Strategies," *Econometrica*, 57, 1171–1214.

Sleesman, Dustin, A. L. G. M. and D. Conlon (2018): "Putting Escalation of Commitment in Context: A Multilevel Review and Analysis," *Academy of Management Annals*, 12, 178–207.

Staw, B. and J. Ross (1980): "Commitment in an experimenting society: An experiment on the attribution of leadership from administrative scenarios," *Journal of Applied Psychology*, 65, 249–260.

Stoneman, P. and G. Battisti (2010): "The Diffusion of New Technology," in *Handbook of the Economics of Innovation*, Elsevier, vol. 2, 733–760.

Taylor, C. R. (1995): "Digging for Golden Carrots: An Analysis of Research Tournaments," *The American Economic Review*, 85, 872–890.

Thomas, C. D. (2021a): "Strategic Experimentation with Congestion," *American Economic Journal: Microeconomics*, 13, 1–82.

——— (2021b): "Strategic Experimentation with Congestion," *American Economic Journal: Microeconomics*, 13, 1–82.

Trajtenberg, M. (1989): "The Welfare Analysis of Product Innovations, with an Application to Computed Tomography Scanners," *Journal of Political Economy*, 97, 444–479.

Tripsas, M. and G. Gavetti (2000): "Capabilities, cognition, and inertia: evidence from digital imaging," *Strategic Management Journal*, 21, 1147–1161.

Van den Steen, E. (2017): "A Formal Theory of Strategy," *Management Science*, 63, 2616–2636.

Verma, Devesh, A. M. and K. K. Sinha (2011): "The Development and Application of a Process Model for R&D Project Management in a High Tech Firm: A Field Study," *Journal of Operations Management*, 29, 462–476.

Wagner, P. A. (2018): "Who goes first? Strategic delay under information asymmetry," *Theoretical Economics*, 13, 341–375.

Whinston, M. D. (2007): "Chapter 36 Antitrust Policy toward Horizontal Mergers," in *Handbook of Industrial Organization*, ed. by M. Armstrong and R. Porter, Elsevier, vol. 3, 2369–2440.

Wong, T.-N. (2018): "Free Riding and Duplication in R & D," *Working Paper*.