THE ROLE OF INFORMATION IN INNOVATION AND COMPETITION

Ufuk Akcigit
University of Chicago

Qingmin Liu
Columbia University

Abstract
Innovation is typically a trial-and-error process. While some research paths lead to the innovation sought, others result in dead ends. Because firms benefit from their competitors working in the wrong direction, they do not reveal their dead-end findings. Time and resources are wasted on projects that other firms have already found to be fruitless. We offer a simple model with two firms and two research lines to study this prevalent problem. We characterize the equilibrium in a decentralized environment that necessarily entails significant efficiency losses due to wasteful dead-end replication and an information externality that leads to an early abandonment of the risky project. We show that different types of firms follow different innovation strategies and create different kinds of welfare losses. In an extension of the core model, we also study a centralized mechanism whereby firms are incentivized to disclose their actions and share their private information in a timely manner. (JEL: O31, D92)

> The failure to report failures means that researchers waste money and effort exploring blind alleys already investigated by other scientists.
>
> "How Science Goes Wrong", *The Economist*, 19 October 2013.

1. Introduction

The cost and benefit of information sharing is not unfamiliar to researchers. In January 2009, Timothy Gowers, a Fields Medalist at Cambridge University, invited

The editor in charge of this paper was Fabrizio Zilibotti.

Acknowledgments: We thank Daron Acemoglu, Alessandro Bonatti, Benjamin Golub, Boyan Jovanovic, Elena Pastorino, Andrew Postlewaite, Andrzej Skrzypacz, Francesco Squintani, Jidong Zhou, the participants at the AEA Winter Meetings, the NSF/CEME Decentralization Conference, and the Fifth Annual Conference on Entrepreneurship and Innovation at Northwestern University for helpful discussions. We thank our editor and four anonymous referees for very helpful comments and suggestions. We also thank Francesc Dilme and Zehao Hu for excellent research assistance. The usual disclaimer applies. Akcigit gratefully acknowledges the financial support of the NBER Innovation Policy Research Grant. The NBER working paper version of our paper (WP 17602) is titled: "The Role of Information in Competitive Experimentation". Akcigit is a Research Fellow at the NBER.

E-mail: uakcigit@uchicago.edu (Akcigit); qingmin.liu@columbia.edu (Liu)

Journal of the European Economic Association August 2016 14(4):828–870

DOI: 10.1111/jeea.12153

all interested mathematicians to openly and jointly tackle a “difficult, unresolved mathematical problem” on his blog. Driven by intellectual curiosity, 27 mathematicians contributed more than 800 mathematical comments, and a generalization, which includes the original problem as a special case, was solved in a mere 37 days.[1] Indeed, the stunning rapidity of this project’s success is that one researcher’s failed ideas and dead-end attempts were not repeated by others, and everyone could focus efforts on the tentatively most promising path. However, intellectual curiosity is not the only motivation for innovation. Incentive schemes, such as patents and prizes, immensely intensify competition in research. In his conquest of the centuries-old Fermat’s Last Theorem, Andrew Wiles worked in complete secrecy for eight years. He even published one of his old papers every six months to keep his colleagues unaware of the direction of his research (Singh 1997).

The two engines of technological progress—competition and innovation—are at odds when it comes to information discovery and sharing. Researchers and firms conduct their research in secrecy and befuddle their competitors; in addition, they do not share information about the exploratory paths that have proven to be fruitless. In private industries, because of the monetary interests involved, the scope of the problem is extravagant. Many firms waste years and millions of dollars on projects that their competitors have already found to be dead ends. In pharmaceuticals, for example, firms in competition to develop a particular drug typically follow similar paths: they try out and then give up on similar compounds due to toxicity or inefficacy. According to a report by the Pharmaceutical Research and Manufacturers of America (P*h*RMA 2011), developing a drug can cost more than $1 billion and take 10–15 years, a significant portion of which arises because firms go through each other’s early failed attempts.[2] Such *dead-end duplications* are common in many sectors with trial-and-error research.

This fact has already alarmed policy makers. For instance, a new project at the Massachusetts Institute of Technology, called New Drug Development Paradigms, is aiming to bring together major drug makers and health authorities to identify and resolve the severe dead-end duplication problem in pharmaceuticals and encourage *precompetitive information sharing* (Singer 2009).[3] While there is general agreement that there should be more information available to competitors about failed research attempts, a better understanding of the economic incentives of competing firms is vital in order to address the question of “precompetitive information sharing” raised by

1. See Nielsen (2012) for more details on Gowers’ Polymath Project.

2. For interested readers, further details on pharmaceutical research and the extent of dead-end duplication can be found in Singer (2009) and P*h*RMA (2011).

3. Another example is the launch by the US Food and Drug Administration and the Critical Path Institute—a nonprofit organization—of a joint program “*Coalition Against Major Diseases*,” which focuses particularly on detailed information sharing about research on Alzheimer’s disease (http://www.c-path.org/camd.cfm). Similarly an international agreement, the “*Bermuda Principles*”, was reached in 1996 to require biologists to share their data on human genome research online (Nielsen 2012). Scientists who refused to share data would receive no grant money. On 14 March 2000, US President Bill Clinton and UK Prime Minister Tony Blair issued a joint statement supporting the Bermuda Principles, asking scientists all around the world to follow these principles (The White House, Office of the Press Secretary, 2000).

policy makers and scientists. How does competition affect firms' research choices and incentives to disclose their findings? Do firms invest too much or too little in risky projects with unknown outcomes and potential dead ends? Which types of firms will most likely pursue risky projects instead of safe projects? What kinds of inefficiencies, if any, arise from the fact that firms can observe only their own failed attempts? Could there be scope for compensating firms that reveal their dead-end findings? Our goal in this paper is to shed light on these important questions.[4]

To study the aforementioned issues, we build a dynamic model of a winner-takes-all research competition between two firms that differ in their arrival rates of innovations. Firms start their competition on a research line that is ex-ante lucrative, but risky—an outcome upon arrival could be *good* or *bad*. A good outcome delivers a one-time lump-sum payoff of Π (e.g., the market value of a drug), while a bad outcome reveals that the research line is a dead end, in which case the payoff is simply 0. In reality, firms have a strong incentive to keep their dead-end findings unknown to their competitors. To capture this incentive, we introduce an additional research line that is ex-ante less lucrative—as aimed for, this structure makes the firm that discovers a dead-end switch secretly to the alternative research. For *tractability*, we assume that the return to this alternative research is low but certain—so we dub the research line "safe". For instance, this alternative safe research corresponds to experiments on the incremental improvement of existing products, while the ex-ante lucrative risky research corresponds to radical innovation whose feasibility is uncertain. We assume that neither the research activity (i.e., which research line the firm is taking) nor the dead-end discovery is publicly observable, while a success is observable (say, through patenting or publication).

Our first contribution is to build a tractable and parsimonious model with the features described previously. Our model features both private outcomes and private actions, as is common in real-world innovation competition. Hence, the analysis of the model requires keeping track of two payoff-relevant beliefs: one about the nature of the risky research and another about the position (research activity) of the competitor. The evolutions of the two beliefs are interdependent. First, to examine the efficiency properties, we solve the model for the case of a single player and then for the case of a social planner who has access to both firms' private information. Our focus is on the decentralized case where both firms compete in a winner-takes-all fashion and the evolution of their private beliefs is key to their strategic choices. We characterize a pure strategy equilibrium in closed form and show that it is unique if the game features enough asymmetry in firms' innovation productivities and payoffs of the research lines. The contrast between the social planner's solution and the decentralized equilibrium outcome is stark and discontinuous in the value of the safe research due to strategic behavior: if the value of the safe research is zero, hiding a dead end on risky research does not bring any strategic advantage. However, if the value of the safe research is strictly positive, however small, the decentralized equilibrium generates

4. Even though our main focus in this paper is on firm-level competition, a similar framework can be used to study the competition among scientists within a firm.

a drastically different equilibrium prediction where firms strictly prefer hiding their dead-end findings. This has severe welfare effects.

It is useful to highlight how the key ingredients in our framework affect the learning dynamics. First, we model two types of outcomes (with a positive prize or a dead end) because such a model is more applicable to the prevalence of trial-and-error types of research competition. Uncertainty about the opponent's discovery type is crucial for our learning dynamics generated by the existence of a dead end. Second, the assumption that the arrival rates on the risky research are independent of the state of the research line implies that there will be no belief updating if research activities and dead-end findings are public. As a result, nontrivial belief updating is entirely driven by the unobservability of dead-end discoveries and private research actions. This is precisely the focus of our analysis. Moreover, this independence assumption implies that efficiency is attainable under perfect information but not otherwise. Hence, the independence assumption isolates and highlights the tradeoff in the applications of our main interest.[5] Third, arrival on the safe line is also stochastic, which affects the learning dynamics indirectly. Upon observing an opponent's discovery on the safe line, a firm can make an inference about the opponent's potential past observations on the risky line, and the extent of this inference in equilibrium turns out to depend crucially on the timing of the safe line discovery. The observational structure in our model is mixed. Actions are not observable unless they lead to a good discovery, but at that point, the competition on that line has ended.

Our second contribution is to identify two major sources of inefficiencies. The first inefficiency arises when one of the firms discovers a dead end and switches silently to the second (safe) research line and the opponent firm still keeps researching on the risky line, even though the competitor had already found it to be a dead end. We call this the *dead-end inefficiency*. We also identify a second inefficiency due to the *information externality*. A firm that has not itself discovered any outcome, nor observed a patent from its rival, could become discouraged about the risky research line and switch to the safe line, even though the risky line is not a dead end, something that never occurs under perfect information. We call this the *information-externality inefficiency.* In addition, we show that when this inefficiency arises, it is always the firm with the lower innovation productivity that switches earlier. While the dead-end inefficiency keeps firms going in a fruitless direction when time and resources should have been used to make discoveries elsewhere (i.e., overinvestment in the wrong project), the information-externality inefficiency prevents firms from concentrating on valuable research (i.e., underinvestment in a valuable project); both effects slow down society's technical progress overall, potentially resulting in a sizable welfare loss. Our numerical analysis suggests that even a very small amount of competition on the safe line generates a very large welfare loss.

Our final contribution is to solve for the required compensation schemes that would incentivize the firms to share their dead-end findings. Asymmetric treatment of winners

5. We discuss the relaxation of this assumption in Section 6.

and losers in the standard patent system creates incentives for research secrecy and concealment of dead-end information. Hence an important lesson we draw from our analysis is that *rewarding failed attempts* is crucial for improving efficiency. Due to the standard difficulties with decentralized information trading (see Arrow 1962 for more on Arrow's Information Paradox), we focus on a third party that ex ante collects monetary installments and rewards the revelation of dead ends as time progresses in an incentive-compatible way. As a result, firms are incentivized to participate in the scheme at any point in time, share their dead-end findings *without any delay* upon their discovery, and follow the first-best decision rules. Notice that while private industries currently reward only profitable positive outcomes, "patents for dead-end discoveries" already exist in some academic professions that publish impossibility results.

Related Literature. Our framework combines elements from several bodies of literature. Innovation as the major source of long-run productivity growth has been the center of a large endogenous growth literature. Monopolistic competition has been the key mechanism in these models.[6] The main premise of these models is that there exists a potential quality or technology improvement, yet the arrival of this improvement is stochastic and affected by the R&D investment of the competing firms. Our paper contributes to this literature by offering a new model in which the existence of a technological improvement is uncertain and therefore firms form their beliefs about the existence of the improvement (or dead end) and update them by observing their competitors' successful findings (patents). Thus, our paper sheds new light on the understanding of the process of innovation.[7]

More directly, our paper contributes to the branch of endogenous growth models with *step-by-step innovations* (see, e.g., Aghion et al. 2001 and Acemoglu and Akcigit 2012). In this class of models, two firms in each sector are engaged in a research competition against each other repeatedly over time. The main feature of these models is that the technology gap between the two competing firms is endogenously determined through the research investments of the leader, follower, or neck-and-neck firms. We follow this literature by introducing a two-stage competition game. In our model, firms start in a neck-and-neck position, and their research investment stochastically determines the technology gap between the competitors. Unlike in that literature, our model features asymmetric information and therefore firms do not observe their competitor but form a belief about it. Moreover, in those models, the technology leader's successful R&D pushes forward the technology frontier and the follower's successful R&D effort typically replicates the steps that were previously already taken by the leader.[8] As a result, the follower's R&D effort is spent on wasteful duplications

6. See the textbooks by Grossman and Helpman (1993), Aghion and Howitt (1998), and Acemoglu (2008) for various models and applications.

7. See also Jovanovic and Rob (1990) and Jovanovic and Nyarko (1996) for a different perspective on the innovation process.

8. One can also consider a model where both the leader's and the follower's innovations push the frontier, as in Acemoglu and Akcigit (2012); in that case, no duplication emerges.

of earlier *successful findings* of the leader. In our model, competing firms not only replicate each other's *dead-end results* as opposed to the successful findings, but also generate an unexpected information externality that leads to the information-externality inefficiency—both types of inefficiencies would have vanished, had private information been made public in our model.

Our analysis shows that some firms might abandon the risky project too early to switch to the safe project. This finding is related to the *discouragement effect* in R&D competitions and technology adoptions, but there are crucial differences. For instance, Fudenberg and Tirole (1985) show that competition and pre-emption motives would make the competing firms adopt technologies too early.[9] Likewise, Katz and Shapiro (1987) show that competition and the threat of being imitated by a competitor could make an innovating firm go after incremental projects. Similarly, step-by-step innovation models (see, for instance, Aghion et al. 2001 or Acemoglu and Akcigit 2012) predict that the risk of imitation makes competing firms invest less in R&D. However, in our model, competition alone is not enough for a firm to abandon the risky project. As explained before, it is the existence of a competitor combined with the "concealment of dead-end information" that makes the weak firm stop the risky project too early.

Our framework studies the dynamics of information sharing with competitors. Therefore it also relates to the literature on interim licensing (see for instance, Spiegel 2008, d'Aspremont et al. 2000, Bhattacharya et al. 1992, and Bhattacharya and Guriev 2006). In those papers, the research outcome is realized in multiple stages and the key focus is identifying the conditions under which interim findings are licensed to competitors. Unlike in those papers, we study the existence of a potential dead-end research path and its welfare implications.

A series of studies has shown that competition among firms and their incentives have important policy implications (see, for instance, Green and Scotchmer 1995, Nickell 1996, Blundell et al. 1999, Scotchmer 2004, and Lerner 2012, among many others). Our paper adds to this discussion by providing a new informational perspective and shows that innovation competition under asymmetric information affects the rival firms differentially, depending on their firm characteristics, which would then have differential effects on welfare and the design of innovation policy.[10]

On the technical side, our paper is related to the strategic bandit literature. Manso (2011) takes an optimal contracting approach to a single-agent experimentation problem and his insight is that an optimal incentive contract involves rewarding failure, though the role of information is not the focus of the model; see also Nanda and Rhodes-Kropf (2012) and the references therein.[11] Strategic experimentation in teams has been studied in the game theory literature; see, for example, Bolton and Harris

9. See Hopenhayn and Squintani (2011) for pre-emption games with private information.

10. Gill (2008) studies patent contests in which a firm can disclose its intermediate research result to signal its commitment to the project.

11. Multi-armed bandit problems have been previously applied to R&D environments. See for example Jensen (1981) and Bhattacharya et al. (1986). See also Pastorino (2011) for a bandit problem with

(1999), Bonatti and Hörner (2011), Heidhues et al. (2014), and Thomas (2011) among many others. The exponential bandit modeling framework has become standard since the seminal contribution of Keller et al. (2005). In these models, there is often one type of outcome, and free-riding rather than competition is a common feature. Early switching is due to the assumption on outcome arrivals that ensure no news is bad news, while in our model, it arises *endogenously* through competition and the concealment of dead-end information. It is important to note that our model would not generate early switching if there were perfect information.

A model that simultaneously features good news and bad news arrivals is particularly suited for analyzing innovations. To our knowledge, our paper is the first in the literature on exponential bandits to simultaneously address learning from competitors (about the risky project) and learning about competitors (what they are working on). From an applied point of view, the paper develops a very flexible tool that delivers ex-post intuitive predictions. The bandit literature usually analyzes fixed games with specific assumptions on the observability of actions and outcomes; we also study efficient information sharing from a mechanism design perspective.[12] Moscarini and Squintani (2010) study R&D competitions when privately informed firms learn from public exit of their opponents. Das (2014) studies competitive bandit problems in which a good risky arm could generate private signals, although the bad risky arm generates no signals; moreover, in his model, actions are publicly observable.

Finally, there is a literature that studies search models of discovery and experimentation. For instance, Fershtman and Rubinstein (1997) investigate a static model of "buried treasures" in which two agents simultaneously rank a finite set of boxes, exactly one of which contains a prize, and subsequently commit to opening the boxes according to that order. There is indeed a dead-end outcome in this model, but due to its static nature, dead-end information is irrelevant and the model does not have a learning element at all. Relatedly, Chatterjee and Evans (2004) offer a dynamic two-arm bandit model of R&D rivalry. In their model, exactly one of the two arms contains a prize but firms do not know which one. In contrast to our central focus here, there is no dead-end discovery in the paper. As a result, searching is always desirable and the issue of dead-end replication does not arise, which is exactly the focus of our paper.

The rest of the paper is structured as follows. Section 2 outlines the model. Section 3 characterizes the equilibrium in a decentralized market. Section 4 provides a numerical example. Section 5 provides an extension to our core analysis and studies a mechanism to incentivize information sharing. Section 6 concludes and also provides a discussion of potential extensions. The Appendix contains the proofs, whereas the

interdependent arms and its applications to R&D and labor markets. The issue of dead-end discovery or replication does not arise in these papers.

12. Bergemann and Hege (1998, 2005) initiate a contracting approach to experimentation problems with applications to venture capital financing and innovation. See also Hörner and Samuelson (2013) and Halac et al. (2013). Halac et al. (2014) study an optimal information scheme in a contest design problem. Bergemann and Välimäki (1996, 2000) study competition and strategic pricing decisions in an environment of experimentation.

Online Appendix provides a number of theoretical extensions of the baseline model and further details of the calibration.

2. Model

Research experimentation is an intrinsically dynamic process. Private outcomes and private actions complicate equilibrium belief formation, especially in the presence of stochastic arrivals on *multiple* research lines. In the following, we attempt to offer the simplest possible dynamic model that captures the essence of the central tradeoffs in such market environments.

2.1. Basic Environment

There are two firms in the economy that engage in research competition in continuous time and maximize their present values with a discount rate $r > 0$. Firms can compete on two alternative research lines: *safe* and *risky*. Each firm can do research on at most one line at a time. For our purpose, we assume that firms start the game with a competition on the risky line.[13] The arrival of outcomes in both lines follows Poisson arrival processes. The safe research is commonly known to deliver a one-time lump-sum payoff $\pi > 0$ upon arrival of an outcome. The risky research has an additional uncertainty besides stochastic arrival. An outcome in the risky research upon arrival could be *good* or *bad*. A good outcome delivers a one-time lump-sum payoff of Π, while a bad outcome reveals that the risky research line is a dead end, in which case the payoff is simply 0. Firms share a common prior $\mu^0 \in (0, 1)$ on the risky research being good.

ASSUMPTION 1. The risky research is ex ante more profitable than the safe research:

$$\mu^0 \Pi > \pi.$$

The two firms differ in their R&D productivities, which are captured in our model by heterogeneous Poisson arrival rates of a discovery. In particular, firm $n \in \{1, 2\}$ has an arrival rate of $\lambda_n > 0$ independent of the research line and has to pay a cost $\lambda_n c > 0$ per unit of time. We assume $\lambda_1 < \lambda_2$. We hence call them weak and strong firms, respectively. We shall write $\Lambda \equiv \lambda_1 + \lambda_2$ as the total arrival rate of both firms.[14]

At time t, a firm can choose one of three options: (1) research on the risky line, (2) research on the safe line, or (3) exit the game with 0 payoff. A firm can change

13. In Online Appendix C, we extend the model by allowing firms to choose simultaneously at $t = 0$ which line to start with (for instance, firms can start with the safe line and switch to the risky line later in the game). This extension complicates the problem, though it is not directly related to our motivation.

14. The only asymmetry between firms is in terms of their arrival rates. Allowing other asymmetries would only complicate the analysis without adding new insights. The role of asymmetry is to rule out coordination equilibria that are not robust. Asymmetry is also a realistic condition from an empirical point of view.

its actions, but it cannot return to the research line it has left. This irreversibility assumption simplifies the analysis of inference/belief-updating without affecting our main focus; it comes at a cost: the calculation of a continuation payoff is more involved.

The firm's research activity is private and unobservable to the public. However, a successful discovery is public.[15] Therefore, a firm is uncertain about which research line its competitor is working on and whether the risky research line has been found to be a dead end, unless it received an arrival on the risky research line or observed a patent by the competitor.

To avoid technical issues associated with continuous-time games, we endow the continuous-time game with two private stages $k = 0, 1$ for each firm in our formal analysis.[16] The game starts at stage 0. In the (common) stage 0, firm n takes the risky research and chooses a stopping time $T_{n,0} \in [0, +\infty]$ at the beginning of this stage. The interpretation is that firm n intends to stay on the risky research line until $T_{n,0}$ *as long as* nothing happens. The game proceeds to stage 1 for firm n at time $t = T_{n,0}$ or when new information arrives at firm n. New information takes one of the following three forms: (i) firm n makes a discovery on the risky research line, (ii) firm n observes a good-outcome discovery from its competitor on the risky line, or (iii) firm n observes a discovery from its competitor on the safe line.

In our game, once an outcome is discovered on a research line, no further positive payoffs will be derived from it. Note that stage 1 is firm n's private stage, because it could be potentially triggered by a private dead-end observation.

If firm n enters its private stage $k = 1$ at $t = T_{n,0}$ when its stopping time expires without observing the arrival of new information, then firm n chooses either "exit" or the "safe research line" with a stopping time $T_{n,1}$. If firm n's private stage $k = 1$ is triggered by the arrival of new information, firm n chooses either "exit" or an available research line together with a stopping time $T_{n,1}$. Note that there is a difference between the two cases. In the latter case, even though new information arrives, firm n can still continue on the risky line if it has not abandoned it yet; in the former case, firm n voluntarily gives up the risky line at $T_{n,0}$ conditional on no arrival of information.

The game for firm n ends if it ever exits, or at $t = T_{n,0} + T_{n,1}$, or if information arrives. Note that the game only consists of at most two private stages for each firm

15. For example, this could be because a patent is needed for a firm to receive the positive lump-sum payoff. Note that in our model, a priori, the incentive for delaying a patent *might* emerge. Strategic patenting will be one of the extensions to our model discussed in Section B.1 of the Online Appendix.

16. This modeling approach was used previously by Murto and Välimäki (2011) in stopping games with public actions. In our model, we allow a firm to react immediately, without a lag, to new information it obtains either by making a discovery on its own or observing potential good discoveries by its opponent. This creates a well-known modeling issue of the timing of events in continuous time. The standard approach adopted in the literature is to focus on Markov strategies that depend only on the beliefs over the risky line, which leads to well-defined outcomes and evolution of beliefs. This approach will not resolve the difficulty in our model with three actions, as a firm's decision depends not only on its assessment of the risky research line, but also on the availability of its outside options in a winner-takes-all competition. For instance, the discovery by the opponent on either research line will not stop the game immediately but obviously affects the continuation game. Moreover, in a multiple-line problem with irreversibility, we need to keep track of the research lines that have been visited in the past (this is not necessary in a one-line problem, as switching research lines ends the game).

TABLE 1. Notation.

Primitives		Values	
π	safe return	w_n^{SS}	firm n's value from competing on the safe line
Π	risky return	w^{SS}	the joint value from cooperating on the safe line
μ^0	prior on the good risky line	w_n^S	firm n's value from monopolizing the safe line
λ_n	firm n 's arrival rate	w_n^R	firm n's value from monopolizing the risky line
Λ	$\lambda_1 + \lambda_2$	w_n^{RR}	the joint value from cooperating on the risky line
c	flow cost per unit of arrival		
Beliefs			
μ_n^t	firm n's beliefs over the risky line at time t		
β_n^t	firm n's period-t belief that the competitor is on the risky line		
b_n^t	firm n's period-t belief that the competitor is on the risky line conditional on the line being bad		

because an observable discovery will remove a research line from the choice set. We focus on a perfect Bayesian equilibrium in pure strategies.

To facilitate the reading of the paper, Table 1 summarizes the notation that appears frequently in the main text.[17]

2.2. The Safe Line

The core of our idea is that competition on the safe research prevents the disclosure of socially efficient information regarding the risky research line. To understand the dynamics of this competition and the effects of the existence of the safe research line, we first shut down the risky research line and consider only the safe research with zero outside options; our analysis here will be used later to determine the equilibrium continuation payoffs in the full game. In the following, we characterize the strategic behavior in three different market structures: monopoly, cooperation, and competition.

2.2.1. Monopoly. Write firm n's monopolistic value from the safe line as w_n^S. Assuming that the firm's strategy is to work on the line until a discovery is made, we can express w_n^S recursively using the following continuous-time Bellman equation:

$$w_n^S = -\lambda_n c dt + e^{-rdt}\left[\lambda_n dt \pi + \left(1 - \lambda_n dt\right) w_n^S\right].$$

17. In choosing this notation, the superscript SS indicates there are *two* firms on the *safe* line; the superscript S indicates that only *one* firm is on the *safe* line. The subscript n indicates that the profit is attributed to firm n.

Here, the first term on the right-hand side (RHS) is the research cost; the second term is the discounted expected instantaneous return, where a lump-sum payoff π is received with an instantaneous probability $\lambda_n dt$; and the third term is the discounted expected continuation payoff.

The Bellman equation immediately gives us

$$w_n^S = \frac{\lambda_n}{\lambda_n + r}(\pi - c). \tag{1}$$

This expression is intuitive. By working on the research line, firm n derives a payoff of $\lambda_n(\pi - c)$ per unit of time (flow payoff), with effective discounting $\lambda_n + r$. From this expression, the firm will research on the safe line if $\pi > c$.

ASSUMPTION 2. $\pi > c$.

It also transpires from the monotonicity of $\lambda_n/(\lambda_n + r)$ in λ_n that the strong firm enjoys larger monopolistic profits.

2.2.2. Cooperation. Next, we consider the cooperative benchmark in which firms maximize their joint value

$$w^{SS} = \frac{\Lambda}{\Lambda + r}(\pi - c),$$

which is positive under Assumption 2. Comparing this with expression (1), the firms now work as one team and hence the arrival rate is $\Lambda = \lambda_1 + \lambda_2$ and the total flow cost is Λc. Since $\Lambda/(\Lambda + r)$ is strictly increasing in Λ, all-firm cooperation is welfare improving over any subset of firms' cooperation, including monopoly as a special case.

2.2.3. Competition. Now consider the winner-takes-all competition between the two firms. Denote firm n's valuation of the safe research line under competition as w_n^{SS}. Assuming that the two firms work on the research line until a discovery is made, the Bellman equation gives us the following intuitive expression:

$$w_n^{SS} = \frac{\lambda_n}{\Lambda + r}(\pi - c). \tag{2}$$

Comparing this with the single-firm case (1), the extra term λ_{-n} in the denominator represents an extra discounting resulting from the competition. Once again, firm n's strategy is optimal if Assumption 2 holds. It is clear that $w_n^{SS} < w_n^S$, meaning that the competition lowers a firm's payoff. Note that $w^{SS} = w_n^{SS} + w_{-n}^{SS}$ is the sum of firms' value under competition. The following proposition summarizes this result.

PROPOSITION 1. *When the research line has a known return, competition is efficient.*

Clearly, the efficiency result depends on the binary action structure. If firms can, in addition, decide on the amount of resources to be allocated to a research line with a convex cost, then competition might induce inefficient resource allocation. This is not the kind of inefficiency we want to identify in this paper. The efficiency benchmark in our model should make this point clear.

3. Equilibrium Analysis of the Model

Now we turn to the full model and analyze dynamic competition with two research lines. We again proceed with three market structures: monopoly, cooperation, and competition.

3.1. Monopoly

If firm n has only the risky research line available, then its monopolistic value can be found using the Bellman equation

$$w_n^R = \frac{\lambda_n}{\lambda_n + r}\left(\mu^0 \Pi - c\right).$$

Note that there is no belief updating in the monopolistic problem. If firm n has only the safe research line available, then similarly its monopolistic value is

$$w_n^S = \frac{\lambda_n}{\lambda_n + r}(\pi - c).$$

Now when the single firm n has two research lines, it will choose when to switch to the safe research line. Firm n's monopolistic value is given by the Bellman equation,

$$v_n = \frac{\lambda_n}{\lambda_n + r}\left(\mu^0 \Pi - c + w_n^S\right) = w_n^R + \frac{\lambda_n}{\lambda_n + r} w_n^S, \tag{3}$$

where $\mu^0 \Pi + w_n^S$ is the expected lump-sum payoff upon an arrival: firm n receives $\mu^0 \Pi$ from the risky research and w_n^S from monopolizing the safe research line. Expression (3) is intuitive. Firm n's expected monopolistic profit from the risky research line is w_n^R, and it also receives the monopolistic profit w_n^S from the safe research line with an arrival rate of λ_n and an effective discount rate of $\lambda_n + r$.

3.2. Cooperation: Planner's Problem

We now consider the case in which firms behave cooperatively to maximize joint value. Two observations are in order, as follows. (i) Firms should share all the information to avoid wasteful research efforts. (ii) Let w^{SS} and w^{RR} be the joint value of the two firms if they work *only* on the safe line and *only* on the risky line, respectively. Using an argument similar to that in the previous section,

$$w^{RR} = \frac{\Lambda}{\Lambda + r}\left(\mu^0 \Pi - c\right) \text{ and } w^{SS} = \frac{\Lambda}{\Lambda + r}(\pi - c).$$

By Assumptions 1 and 2, we have $w^{RR} > w^{SS} > 0$.

The planner's strategy space is larger than the monopolist's problem. In particular, the problem involves the optimal allocation of joint efforts. Therefore, a more interesting question is how to allocate the joint efforts and, in particular, whether splitting the research lines between the two firms is more desirable. We shall show that the first best allocation of efforts requires that both firms work on the risky line until a discovery is made (which is made public immediately) and then both switch to the safe line. Splitting the task is never optimal.

PROPOSITION 2. *Under Assumptions 1 and 2, the strategy that maximizes joint value is for both firms to work on the risky line together until a discovery is made, and then both switch to the safe line. The joint value is given by*

$$V = w^{RR} + \frac{\Lambda}{\Lambda + r} w^{SS}, \tag{4}$$

and if firm n is awarded the good discovery it makes, then its value is

$$V_n = \frac{\lambda_n}{\Lambda} w^{RR} + \frac{\Lambda}{\Lambda + r} w_n^{SS}. \tag{5}$$

Proof. See the Appendix. □

The interpretation of the joint value under this strategy is as follows. Recall that w^{RR} is the joint value of researching only on the risky research until an outcome is found. When the firms follow a strategy of researching on the risky line and then switching to the safe line upon discovery, this also adds the continuation value of the safe research on top of w^{RR}. A discovery on the risky line arrives at the rate Λ and the firm's continuation payoff from the safe research upon arrival is simply w^{SS}.

Proposition 1 and Proposition 2 together imply that absent either risky innovation or market competition, the R&D game has an efficient outcome. Next, we show that the interaction of risky innovation and competition leads to undesired inefficiencies.

3.3. Competition in a Decentralized Market

When it comes to competition, which research line a firm is working on is private information and only the good discovery is observable. We shall now demonstrate how learning and private beliefs become tractable in our model.

3.3.1. Learning and Private Beliefs. Write μ_n^t as firm n's private belief that the risky research line contains a good outcome at time t (which obviously depends on the realization of private and public histories). Write β_n^t as the probability that firm n assigns to its opponent, firm $-n$, being on the risky line at time t. Denote by b_n^t the probability that firm n assigns to its opponent being on the risky line at time t conditional on the fact that the risky line is bad.

Suppose both firms start on the risky line and switch only upon an observation. If firm n does not observe anything—neither from itself nor from its opponent—from t to $t + dt$, firm n will update μ_n^t using Bayes' rule as follows:

$$\begin{aligned}\mu_n^{t+dt} &= \frac{\mu_n^t \left(1-\lambda_{-n}dt\right)\left(1-\lambda_n dt\right)}{\mu_n^t \left(1-\lambda_{-n}dt\right)\left(1-\lambda_n dt\right) + (1-\mu_n^t)\left[1-(1-b_n^t)\lambda_{-n}dt\right]\left(1-\lambda_n dt\right)} \\ &= \frac{\mu_n^t \left(1-\lambda_{-n}dt\right)}{\mu_n^t \left(1-\lambda_{-n}dt\right) + (1-\mu_n^t)\left[1-(1-b_n^t)\lambda_{-n}dt\right]}.\end{aligned}$$

Note that the final expression is independent of $(1 - \lambda_n dt)$; that is to say, firm n does not learn from the fact that it does not observe anything from its own research. This is because the arrival rate λ_n is independent of the type of the outcome (see the previous discussion).

The interpretation of the previous equality is as follows. The numerator measures the probability that the opponent does not make a (public) discovery and the risky line is good. The denominator measures the probability that firm n does not observe anything from its opponent—when the risky research is a dead end, the only observable discovery from its opponent is on the safe line, which occurs with probability $(1 - b_n^t)\lambda_{-n}dt$, and hence the probability of observing nothing from $-n$ is $1 - (1 - b_n^t)\lambda_{-n}dt$.

From the previous Bayesian updating, we derive the law of motion for private beliefs:[18]

$$\dot{\mu}_n^t = -\mu_n^t \left(1-\mu_n^t\right) b_n^t \lambda_{-n}. \tag{6}$$

The critical feature of the learning is that when the opponent makes a discovery faster (i.e., when λ_{-n} is larger), then firm n learns faster. The intuition is as follows. As λ_{-n} increases, the opponent will discover an outcome on the risky research sooner. Therefore, if no good outcome is observed from the opponent over a fixed period of time, it is more likely that the opponent actually found a dead end. Therefore, everything else equal, *the weak firm becomes more pessimistic than the strong firm* on the risky research over time with no discovery.

If firm n knows that a bad (dead-end) outcome has arrived before t, then $\mu_n^t = 0$; if n knows that the good outcome has occurred before t, then $\mu_n^t = 1$.

Learning with Stopping Strategies. Suppose both firms work on the risky line before $T > 0$ until a discovery is made. How will the private beliefs evolve? First, at any $t \leq T$, if firm n has not observed anything from its opponent or from its own research, then

$$\beta_n^t = \frac{e^{-\lambda_{-n}t}}{e^{-\lambda_{-n}t} + \left(1-\mu^0\right)\lambda_{-n}te^{-\lambda_{-n}t}} = \frac{1}{1+\left(1-\mu^0\right)\lambda_{-n}t}. \tag{7}$$

18. To see this, subtract μ_i^t from both sides of Bayes' formula, divide them by dt and then take the limits.

We need to interpret this formula: $e^{-\lambda_{-n}t}$ is the probability that the opponent firm $-n$ does not make any discovery by time t; $(1-\mu^0)\lambda_{-n}te^{-\lambda_{-n}t}$ is the probability that the opponent makes one dead-end discovery and that is the only discovery by time t; since the arrival rate is λ_{-i}, the probability of *one and only one* arrival by time t is $\lambda_{-i}te^{-\lambda_{-i}t}$.[19] The denominator in (7) is the total probability of no observation from the opponent, which consists of two pieces: the probability of no arrival, $e^{-\lambda_{-n}t}$, and the probability of only one private (dead-end) arrival $(1-\mu^0)\lambda_{-n}te^{-\lambda_{-n}t}$. The opponent will stay on the risky line only when there is no arrival by $t \leq T$. This is reflected in the numerator of (7).

Similarly, if firm n has not observed anything from its opponent and from its own research, then conditional on the risky research having a dead end,

$$b_n^t = \frac{e^{-\lambda_{-n}t}}{e^{-\lambda_{-n}t} + \lambda_{-n}te^{-\lambda_{-n}t}} = \frac{1}{1+\lambda_{-n}t}. \tag{8}$$

Note that b_n^t is conditional on the risky research having a dead end, and hence, $(1-\mu^0)$ is excluded from Bayes' formula (7). Substituting equation (8) into the filtering equation (6), we obtain

$$\dot{\mu}_n^t = -\mu_n^t\left(1-\mu_n^t\right)\frac{\lambda_{-n}}{1+\lambda_{-n}t}. \tag{9}$$

As this formula demonstrates, even though the rate of discovery λ_{-n} is constant over time in our model, the rate of learning from no observation, $(\lambda_{-n})/(1+\lambda_{-n}t)$, changes hyperbolically in time. The following lemma provides the explicit form for the belief.

LEMMA 1. *Under the stopping strategies described previously, the belief of firm n at time $t \leq T$ that the risky research has a good outcome is*

$$\mu_n^t = \frac{\mu^0}{1+\left(1-\mu^0\right)\lambda_{-n}t}. \tag{10}$$

Proof. See the Appendix. □

Now, consider the case in which firm n has not discovered anything from its own research but observes the opponent's discovery on the safe research at $t \leq T$. Given the stopping strategy that firm $-n$ adopts, firm n could infer that the opponent has already discovered a dead end on the risky research previously and has since switched to the safe research. Therefore, in this case, $\mu_n^t = 0$.

Next, consider the case in which firm n has not discovered anything through its own research but observes the opponent's discovery on the safe research at $t > T$. Then

19. $\int_0^t e^{-\lambda_{-i}s}\lambda_{-i}e^{-\lambda_{-i}(t-s)}ds = \lambda_{-i}te^{-\lambda_{-i}t}$.

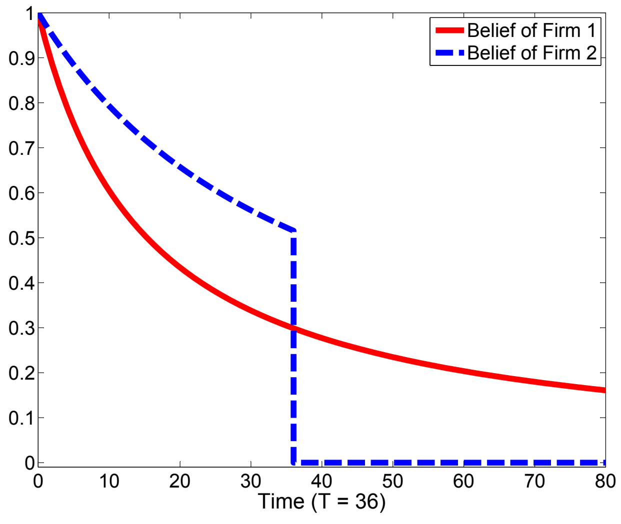

FIGURE 1. b_n^t: beliefs on the opponent being on the risky arm.

there is no updating $\mu_n^t = \mu_n^T$, and in fact, this observation is valid as long as firm $-n$ switches at time T, and it does not matter when firm n switches. This observation is immediate from the following:

$$\mu_n^t = \frac{\mu_n^T e^{-\Lambda(t-T)}}{\mu_n^T e^{-\Lambda(t-T)} + \left(1-\mu_n^T\right) e^{-\Lambda(t-T)}} = \mu_n^T.$$

Finally, if firm n has not discovered anything from its own research at $t > T$, its belief μ_n^t is still μ_n^T. Note that there is a very interesting discontinuity: when firm $-n$ makes a discovery on the safe research at or before T, then μ_n^t jumps down to 0, while, if the discovery is made right after T, the belief is constant at μ_n^T, as if nothing had occurred. This discontinuity illustrates the intricacy of the belief updating process and strategic incentives in our model.

With the previous discussion as a precursor, Figures 1 and 2 depict the evolution of beliefs, b_n^t and μ_n^t, conditional on no arrival under the following pair of stopping strategies: until an observation reveals the nature of the risky line, firm 1 stays on the risky line until $T > 0$, and firm 2 sticks to the risky line.[20]

Of course, a priori, there is no guarantee that the *equilibrium* evolution of beliefs will be as clean as conjectured previously. We confirm this in the next section.

3.3.2. Equilibrium. Recall that we assume firm 1 is weaker than firm 2 in the sense that $\lambda_1 < \lambda_2$.

20. The parameters come from a simple numerical exercise provided in Section 4.

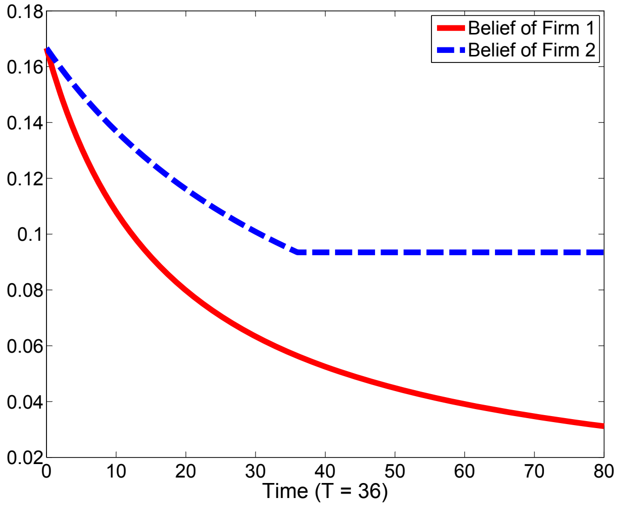

FIGURE 2. μ_n^t: beliefs on the goodness of the arm.

PROPOSITION 3. *Under Assumptions 1 and 2, there is a pure strategy perfect Bayesian equilibrium in which both firms start on the risky research and switch silently to the safe line upon a dead-end discovery. In this equilibrium,*

- *unless an outcome is observed, the strong firm will not stop, and the weak firm (firm 1) will switch to the safe research line at*

$$T = \frac{1}{(1-\mu^0)\,\lambda_2}\left[\frac{\mu^0\Pi}{\pi}\frac{r+\Lambda}{r+\Lambda-\lambda_1\left(\frac{\pi-c}{\pi}\right)}-1\right]+\frac{\lambda_1\frac{\pi-c}{\pi}}{(r+\Lambda)\left(r+\Lambda-\lambda_1\frac{\pi-c}{\pi}\right)},$$

- *if the first news that a firm observes from its opponent before T is a good outcome of the risky research, then both firms switch to the safe research,*
- *if the first news that a firm observes from its opponent before T is an outcome on the safe research, then both firms exit,*
- *if firm 2 observes a good outcome on the risky research after T, it will switch to the safe research if it is still available.*

Finally, if there is enough asymmetry across research lines and players, namely if $\mu^0\Pi/\pi$ and λ_2/λ_1 are large enough, then the previous conditions describe the unique pure strategy equilibrium outcome.

Proof. See the Appendix. □

In contrast to the planner's problem, this result shows that a small prize on the safe research changes the incentives of competing firms discontinuously and distorts the market outcome. In this equilibrium, the weak firm abandons the risky research too early compared to the first-best scenario in which both firms stay on the risky research until a discovery is made. Indeed, this is the case even when λ_1 approaches λ_2. This equilibrium also reveals that the two asymmetric firms generate different types of inefficiencies absent from a discovery on the safe line. First, the strong firm generates wasteful duplicative R&D from the time that the weak firm discovers a failure until it discovers the failure itself or the weak firm discovers the safe line before T. Second, the weak firm generates wasteful R&D only from the time that the strong firm discovers a failure until its switching time T or the time at which the strong firm discovers the safe line. Moreover, the weak firm generates inefficiency from the time it switches until the strong firm discovers an outcome in the risky line, due to early switching. In short, the weak firm endures two kinds of inefficiencies: *information-externality and dead-end inefficiencies*, while the larger firm endures only the *dead-end inefficiency*. We offer a more detailed analysis of welfare effects, as well as other effects of decentralized competition, via a numerical example in the next section.

We also want to comment on the role of asymmetry. If firms are symmetric or payoffs on both research lines are close, mixed strategy equilibria are also possible. In this case, the coordination failures would generate the possibility of abandonment of a viable research line.[21]

The following proposition provides a comparative statics analysis with respect to the parameters of the model.

PROPOSITION 4. *The equilibrium stopping time T is increasing in μ^0 and Π, and decreasing in λ_2 and π.*

Proof. See the Appendix. □

These comparative statics are intuitive. As μ^0 and Π become larger and π becomes smaller, the risky line becomes more attractive. However, when λ_2 becomes larger, the weak firm updates its belief downwards faster. The response of T with respect to λ_1 is nonmonotonic as it affects both the weak firm's payoffs in both lines simultaneously.

4. A Numerical Example

In this section, we provide a numerical analysis, taking pharmaceutical research competition as an example. Our goal is to illustrate the behavior and welfare implications of the model and highlight its general quantitative features for reasonable parameter values. Our model has seven parameters: r, μ^0, Π, π, c, λ_1, and λ_2. Our strategy is to calibrate the model to the clinical trial stage of the pharmaceutical research during the late 1990s as these are the years for which we have information both on the cost of drugs and on the profits of the companies. Table 2 summarizes the parameter values in our example.

21. Please see the Online Appendix for additional details.

TABLE 2. Parameter values (monthly) and equilibrium stopping time.

r	μ^0	λ_1	λ_2	c	Π	π	T
0.4%	17%	2.6%	6.5%	\$63 million	\$1.4 billion	\$87 million	36 months

TABLE 3. Comparison of decentralized and planner's solutions.

Moment	Decentralized	Planner's
Average time to develop a risky drug	14.9 years	11 years
Average cost to develop a risky drug	\$499 million	\$382 million
Fraction of risky drugs invented by firm 1	28%	29%
Average risky experimentation by firm 1	16.1 months	10.9 months
Average risky experimentation by firm 2	13.8 months	10.9 months
Average safe experimentation by firm 1	9.1 months	10.9 months
Average safe experimentation by firm 2	11.7 months	10.9 months
Average wasteful risky research investment by firm 1	9.6 months	0
Average wasteful risky research investment by firm 2	11.4 months	0

These parameters come from a simple calibration exercise in which we rely on reports by the Pharmaceutical Research and Manufacturers of America P*h*RMA (2011). The details of the parameter choices are described in Online Appendix D.

4.1. Summary Statistics

Table 3 summarizes the key variables given the parameters in Table 2. Each firm n starts on the risky line with an initial belief $\mu_n^0 = 1/6$. As time elapses, firms receive outcomes according to the Poisson process. Note that firm 2 observes an outcome roughly 2.5 times more frequently than the weak firm 1 (λ_2/λ_1). Because firm 2 receives an outcome faster, its average experimentation time on the risky line is shorter by around 13.8 months as opposed to 16.1 months for firm 1. Note that this is despite the fact that firm 1 follows a cutoff rule according to which it switches to the safe line at $T = 36$ if it does not observe an outcome either from itself or from its competitor.

The associated beliefs under this strategy are depicted in Figures 1 and 2.

Figure 3 depicts the distribution for experimentation durations on the risky line in each trial. The first point to note is the spike at $t = 35$. In almost 12% of the trials, firm 1 does not observe any outcome and follows its equilibrium cutoff strategy, switching to the safe line at $t = T$. Second, compared to firm 1, firm 2's distribution has more mass at lower durations. This is due to the fact that firm 2 has a faster arrival rate, which allows it to discover the true nature of the risky line more quickly. Finally, in the planner's economy, information sharing increases the effective arrival rate for both firms ($\lambda_1 + \lambda_2$). This shifts the distribution of experimentation durations to the left and

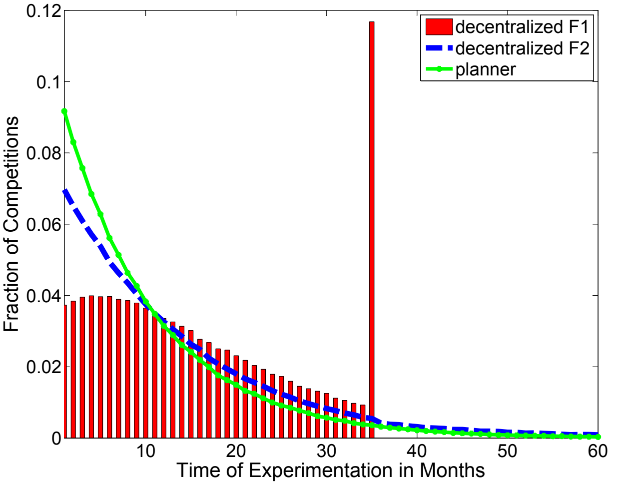

FIGURE 3. Duration of risky experimentation.

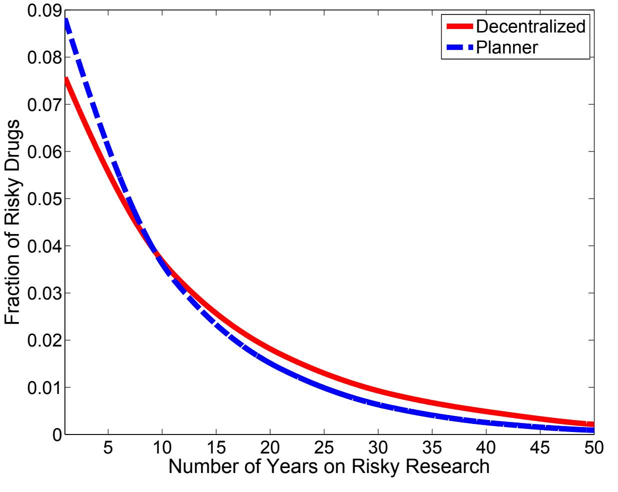

FIGURE 4. Total firm years spent until the next successful drug.

hence reduces the average time spent on the risky line to 10.9 months, which is 32% and 21% lower than the average experimentation times for firms 1 and 2, respectively.

Next, we study the time that firms spend on risky research between two consecutive risky drug inventions. Figure 4 plots the results of the numerical simulations. In the decentralized economy in which firms have private information about their R&D outcomes, firms spend on average 14.9 years on the risky line per drug. Note that some of this time is spent on research in a line that the competitor already knows is a dead end. The planner's economy avoids this problem, and firms spend eleven years—that is 26% less time—on the risky line per drug.

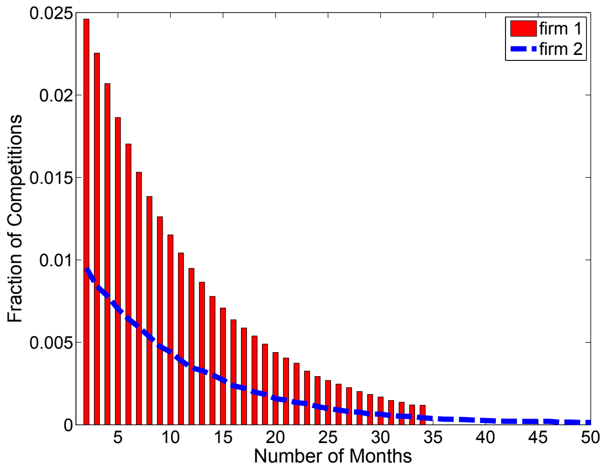

FIGURE 5. Dead-end replications (dead-end inefficiency).

It is also important to understand the sources of inefficiencies in the economy. The decentralized economy differs from the planner's economy in two major dimensions. First, when a firm discovers a dead end on the risky line before T, it switches to the safe line without sharing this information with the competitor. As a result, the competitor is wasting R&D dollars on a research line that is already known to be a dead end. This is what we call the *dead-end inefficiency.* Figure 5 plots the distribution of the number of periods spent on research in a dead end. Note that the maximum wasteful R&D by firm 1 has an upper bound of T, due to the cutoff strategy, which mitigates the welfare loss (however, as will be shown in what follows, this strategy increases the second type of inefficiency). Because firm 2 learns the true nature of the line faster, firm 1 spends more time on a dead-end risky line before T. However, while firm 2 incurs wasteful R&D spending less frequently before T, it is the only firm that can potentially stay longer on a dead-end research line. The average dead-end replication time is 9.6 months for firm 1 and 11.4 months for firm 2.

Figure 6 describes the second source of inefficiency: *information externality.* The planner prefers both firms to experiment until an outcome is found on the risky line. However, in the decentralized economy in which firms do not observe the private information of their competitors, they become pessimistic about the outcome on the risky line, as time elapses. Firm 1 switches to the safe line at time T even in situations where firm 2 has not received any information about the risky line by then. This generates missing experimentations by firm 1 due to early switching, which are plotted in Figure 6.

Finally, we illustrate the monetary cost of the problem in Figure 7, which plots the distribution of the total amount of R&D dollars spent between two consecutive risky drugs. In the decentralized economy, firms spend on average \$499 million on a risky drug, a significant portion of which is wasted due to the two aforementioned

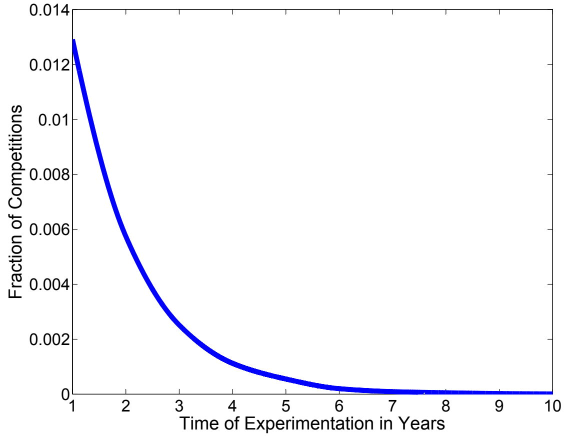

FIGURE 6. Missing experimentations (information-externality inefficiency).

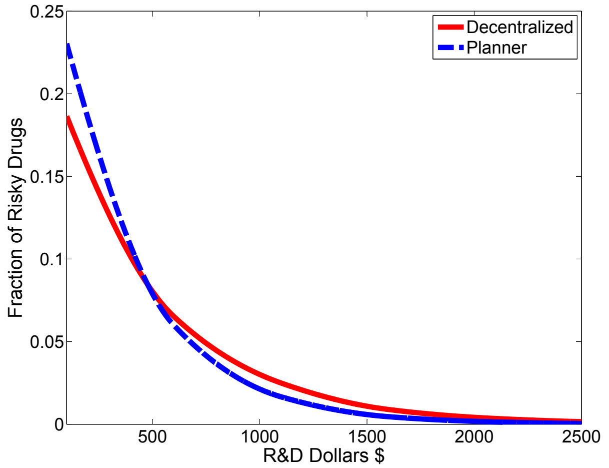

FIGURE 7. Distribution of R&D dollars per drug.

inefficiencies. Firms spend on average \$382 million in the planner's economy, which is 23% less.

The following section discusses the sources of these inefficiencies in greater detail.

4.2. Two Types of Inefficiencies: Dead End and Information Externality

In this section, we focus on two different types of inefficiencies demonstrated in our equilibrium. We consider three regimes: the first-best regime (FB) is the cooperation setup with information sharing, the decentralization regime (D) is the decentralized market without information sharing, and the intermediate regime (I) has full information sharing, but artificially requires the weak firm 1 to stop at T, the stopping time in regime D.

Let us denote the welfare associated with the regime α as W_α, where $\alpha \in \{FB, D, I\}$. Therefore, $W_{FB} - W_I$ is the welfare loss due to early switching only (excluding the information externality upon the discovery of bad news), and $W_I - W_D$ is the welfare loss due to the information externality—socially efficient information of a dead-end finding is not disclosed.

From Proposition 2, we know that

$$W_{FB} = w^{RR} + \frac{\Lambda}{\Lambda + r} w^{SS}.$$

Because the intermediate regime differs from the first-best regime only after T, we have

$$W_{FB} - W_I = \lambda_1 \left[\left(\mu_0 \Pi + w^{SS} \right) - \left(\pi + w_2^R \right) \right] \frac{e^{-(\Lambda + r)T}}{\Lambda + r},$$

where $\lambda_1(\mu_0 \Pi + w^{SS})$ and $\lambda_1(\pi + w_2^R)$ are firm 1's contribution to the total welfare (measured in flow payoffs) when firm 1 works on the risky line and the safe line, respectively; $e^{-\Lambda T}$ is the probability that a discovery has not been made on the risky research by T.

Finally, note that the difference between regime (I) and regime (D) arises only when the risky research is a dead end. In this case, a dead-end discovery is not observable to the opponent, unless a subsequent discovery on the safe line is reported before T. Therefore, we need again to consider the probability that only one discovery is made by the same firm n before t, which is given by Pr(one arrival before t) $= \lambda_n t e^{-\lambda_n t}$. Using this fact, we obtain[22]

$$W_I - W_D = \left(1 - \mu^0\right) \frac{\lambda_1 \lambda_2}{r + \Lambda} \left[2 \frac{\pi}{r + \Lambda} \left(1 - e^{-(r+\Lambda)T} \right) \right.$$
$$\left. - T e^{-(r+\Lambda)T} \left(\pi - \frac{\lambda_1 c}{r + \lambda_2} \right) \right].$$

Table 4 summarizes the numerics. Note that firms do not want to share the dead-end discovery on the risky line because of the competition on the safe line, which has a per unit of arrival rate net return $\pi - c$.

The finding is striking. We notice that even if the net return on the safe line is only \$1, the incentive of preventing the opponent from competing for this \$1 causes a total efficiency loss of \$19.3 million, which amounts to 12% of the first-best welfare level! The logic, as we have already pointed out, is that this \$1 completely changes

22. This follows from

$$W_I - W_D = (1 - \mu^0) \left\{ \begin{array}{l} \int_0^T \lambda_2 t e^{-\lambda_2 t} e^{-\left(r+\lambda_1\right)t} \lambda_1 \pi dt + \int_0^T \lambda_1 t e^{-\lambda_1 t} e^{-\left(r+\lambda_2\right)t} \lambda_2 \pi dt \\ + \int_T^\infty \lambda_1 T e^{-\lambda_1 T} e^{-\left(r+\lambda_2\right)t} \left[e^{-\lambda_1 (t-T)} \lambda_2 \pi + \left(1 - e^{-\lambda_1 (t-T)}\right) \lambda_2 c \right] dt \end{array} \right\}.$$

TABLE 4. Welfare analysis.

$\pi - c$	$W_{FB} - W_I$	$W_I - W_D$	$W_{FB} - W_D$	W_{FB}	$\frac{W_{FB}-W_D}{W_{FB}}$
Level of Competition	Information-externality Inefficiency	Dead-end Inefficiency	Total Inefficiency	First-best Welfare	Percentage Inefficiency Loss
\$1	\$0.024 m	\$19.3 m	\$19.3 m	\$162.9 m	12%
\$1 m	\$0.026 m	\$19.5 m	\$19.6 m	\$163.8 m	12%
\$10 m	\$0.044 m	\$21.9 m	\$22.0 m	\$172.0 m	13%
\$30 m	\$0.364 m	\$31.7 m	\$32.1 m	\$254.5 m	13%

the incentives to share private information. Without it, the firm does not lose anything from information sharing.

REMARK 1. Note that the dead-end inefficiency is much larger than the information-externality inefficiency. We should *not* be optimistic about the information-externality inefficiency. Indeed, early switching due to the information externality delays the discovery on the risky line by almost four years for the same set of parameters as we demonstrated previously. If consumers' welfare is taken into account, then information-externality will have a much larger implication.

5. Extension: Incentivizing Information Sharing

In this section, we shall consider an extension to our core analysis and explore the possibility of a mechanism that incentivizes information sharing. It should be emphasized that we do not suggest that our mechanism is practical, because, as in the theoretical mechanism design literature, our mechanism depends on the details of the model; rather, we want to investigate theoretically the outreach and the limits of the simple idea of trading dead-end discoveries. The idea is to create a *centralized institution* to reward dead-end discoveries. This is the counterpart of the prevailing practice of rewarding good-end discoveries through patents and prizes. After all, many professions publish and reward dead-end discoveries and impossibility results. We focus on the case where outcomes are verifiable. Similar to good-outcome patenting where firms prove that their experiments lead to the solution of a problem (e.g., a drug curing a disease), we assume that firms can provide their research results and data to prove their dead-end findings (similar to the data policy of academic journals and proofs of impossibility results).

REMARK 2. One important question to answer is why there is a need for a mechanism designer instead of allowing firms to trade dead-end discoveries in a decentralized market or to sign contracts among themselves. This is the core of the classic problem of information trading, as pointed out by Arrow (1962) in an argument for patenting through centralized institutions. Information is different from standard commodities. The buyer of information, once the buyer learns the information or verifies it, obtains what he needed in the first place and no longer has incentive to pay more. This problem discourages information trading in a decentralized market. Therefore, a mediator is often necessary for the sale of information.

5.1. Feasible Mechanisms

The mechanism must be dynamic in nature to accommodate the stochastic arrival. Ideally, a dynamic mechanism that enforces information disclosure should satisfy the following properties: (i) budget balance; (ii) a firm at *any point in time* should be allowed to walk away from the mechanism (i.e., we face a design problem in which firms cannot commit to their future actions); (iii) a firm should not walk away from the mechanism at some point and then come back in the future to take advantage of the information accumulated during its leave; and (iv) a dead-end outcome should be made public immediately upon its discovery with no delay.

One particular issue with this type of mechanism is that if a firm walks away (off the equilibrium path), the other firm is left wondering what the firm has actually observed that made it leave; there is a myriad of off-path beliefs, and each belief can potentially support a different decentralized continuation equilibrium play. Thus, the parameters of the mechanism will depend on the specification of off-path beliefs. Note, however, that this issue must emerge in any dynamic mechanism design problem where agents could receive new information over time when agents cannot commit to their plan of action at time 0.

The off-path beliefs have to be realistic and robust to perturbations. Indeed, we could think of perturbation of firm strategies in the game-theoretic tradition of trembling-hand perfection, or alternatively, we can think of a rare, random exogenous shock that forces a firm to leave the mechanism. In the latter case, exiting the mechanism becomes an on-path behavior and beliefs follow directly from standard Bayes' updating. These considerations lead us to adopt the following specification of off-path beliefs.

- If a firm quits the mechanism at some point, which is off the equilibrium path, then the other firm's belief does not suddenly change.

We shall design a mechanism with these properties. The mechanism simply states the following. At any time t, each firm can report a failure it discovered to a mediator; if firm n reports a failure, then firm $-n$ will be liable to pay p_n^t to firm n, and the mechanism concludes. For example, firm n can deposit p_n^t in a neutral account at time t managed by the mediator. Our goal is to find the range of p_n^t that satisfies the incentive conditions.

Henceforth we shall restrict our attention to a constant price path such that $p_n^t = p_n$. The following lemma summarizes the conditions on the prices.

LEMMA 2. *Consider a set of prices p_1, p_2 that satisfy no-delay, no-walk-away, and participation constraints. Then these constraints can be formally described as follows.*

1. No-delay condition:

$$\frac{\lambda_{-n}\lambda_n}{\Lambda + r}(\pi - c) \leq \left(\lambda_{-n} + r\right) p_n + \lambda_{-n} p_{-n}, \text{ for } n = 1, 2. \tag{11}$$

2. No-walk-away with a dead end:

$$p_1 \geq \frac{\lambda_1 \lambda_2}{(\Lambda + r)^2}(\pi - c) \text{ and } p_2 \geq \left[1 - e^{-(\Lambda + r)T}\right] \frac{\lambda_1 \lambda_2}{(\Lambda + r)^2}(\pi - c). \tag{12}$$

3. Participation constraint:

$$\underline{K} \leq \lambda_1 p_1 - \lambda_2 p_2 \leq \bar{K}. \tag{13}$$

Proof. See the Appendix. □

Now we are ready to state the main proposition of this section.

PROPOSITION 5. *Each price vector (p_1, p_2) that satisfies conditions (11) and (13) characterizes a mechanism that restores efficiency: both firms work on the risky research until a discovery is made and then switch to the safe research; firm n reports a dead-end discovery immediately upon its discovery and receives a payment p_n from its competitor.*

Proof. Note that the set of price vectors (p_1, p_2) that satisfy (11)–(12) is nonempty. Indeed, we can set $p_1 = (\lambda_2 p_2 + \bar{K})/\lambda_1$, which satisfies (13). By setting p_2 large enough, all other constraints will be satisfied simultaneously. By definition, firms share their information without delay under the mechanism with (p_1, p_2). The result then follows. □

There is a continuum of price vectors that satisfy conditions (11)–(12). One way to refine this set of price vectors is to introduce a liability constraint. Instead of pushing in this direction, we characterize the "cheapest" prices that are enough to restore efficiency. To do this, we minimize the flow transfer $\lambda_1 p_1 + \lambda_2 p_2$ over all mechanisms.

5.2. *Minimum Implementable Transfers*

Formally, minimizing the flow transfer $\lambda_1 p_1 + \lambda_2 p_2$ over all mechanisms is the following linear programming problem:

$$\min_{(p_1,p_2)} \quad \{\lambda_1 p_1 + \lambda_2 p_2\}$$

$$\text{subj. to:} \quad \text{C1:} \frac{\lambda_1\lambda_2}{\Lambda + r}(\pi - c) \le (\lambda_1 + r)\, p_2 + \lambda_1 p_1,$$

$$\text{C2:} \frac{\lambda_1\lambda_2}{\Lambda + r}(\pi - c) \le (\lambda_2 + r)\, p_1 + \lambda_2 p_2,$$

$$\text{C3:} \frac{\lambda_1\lambda_2}{(\Lambda + r)^2}(\pi - c) \le p_1,$$

$$\text{C4:} \left[1 - e^{-(\Lambda+r)T}\right] \frac{\lambda_1\lambda_2}{(\Lambda + r)^2}(\pi - c) \le p_2,$$

$$\text{C5:} \underline{K} \le \lambda_1 p_1 - \lambda_2 p_2 \le \bar{K}.$$

The set of binding constraints in this program is determined by primitive parameter values of $c, \lambda_n, r, \pi, \mu_0$, and Π. We present numerical solutions using the previous set of parameters. The interesting finding is that the cost of the mechanism is quite minimal relative to the size of the recovered welfare loss.

In the numerical computations, the two binding constraints of the mechanism are the no-delay condition for firm 1, (C1) and the no-walk-away condition for firm 2 (C4). Figure 8 plots the prices dictated by the minimum transfer mechanism as a function of the competition level on the safe research line.

Two features stand out in Figure 8. First, the price that each firm has to pay to compensate its competitor is increasing in the level of the competition on the safe research line. Second, the price that firm 1 receives (p_1) is always higher than that of firm 2, since sharing information on a dead-end finding means that both firms will now compete on the safe line. For firm 1, this entails a larger reduction in value because it will then face a stronger competitor (firm 2).

6. Concluding Discussion and Future Research

The goal of this paper has been to uncover the potential inefficiencies in research competitions due to dead-end replication. We offered a parsimonious two-line research competition model with two asymmetric firms. We identified two types of inefficiencies that arise in this model and showed that different firms incur different types of inefficiencies. The efficiency loss is significant, and we have discussed a simple mechanism to improve efficiency. We have made several simplifying assumptions to highlight the effects of a dead-end discovery and asymmetric information.

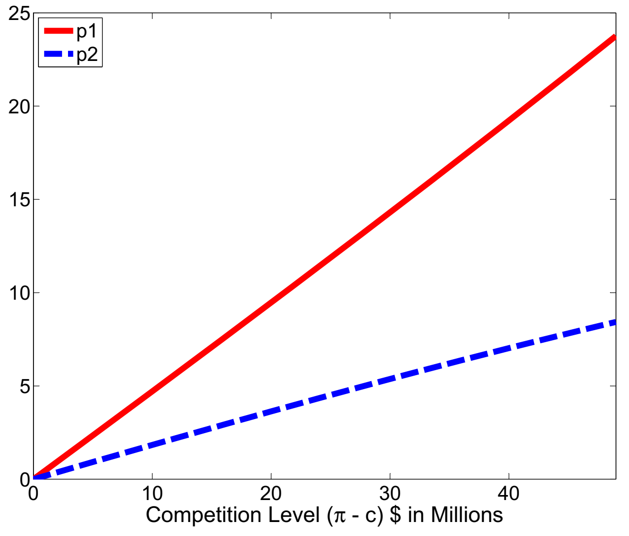

FIGURE 8. Competition versus prices in min price mechanism.

One final point regarding the robustness of our analysis is worth mentioning. In our framework, firms receive a lump-sum payoff from a good discovery immediately. Could an endogenous patenting decision alter the results? More specifically, could there be a gain from strategically delaying the reporting of the good outcome? As we show in Section B.1 of the Online Appendix, the answer is no. In our framework, the cost of potentially losing the valuable successful innovation in the risky arm is so high that the gain from hiding it is always exceeded by its cost. Our analysis has focused on two firms. Indeed, it will be natural to extend the analysis to multiple firms. See Section B.2 of the Online Appendix for additional discussions.

There are interesting possible extensions of our framework for future research. One possible extension is on state-dependent arrivals of information. In particular, one might allow the arrival rate of the research outcome in the risky arm to be a function of the state of the arm (for instance, the researcher might learn the nature of the arm only if it is good). A source of exogenous learning would show up in this environment. As we discuss further in Section B.3 of the Online Appendix, the endogenous learning in our model and this potential exogenous learning have different dynamics and it would be a fruitful direction to study the interaction between these two channels.

Another important direction would be to study the macroeconomic implications of wasteful dead-end duplications. Macro data on innovation and R&D spending in the United States exhibit a worrisome time-series pattern: the ratio of registered innovation counts to total innovation efforts in the United States has been steadily decreasing over time. This pattern is even more severe in the pharmaceutical sector. These patterns are discussed in more detail in Section B.4 of the Online Appendix. Through the lens of our framework, one possible explanation would be the rise in dead-end duplications. More specifically, a rise in the value of innovation or changes in the beliefs on the existence of a successful outcome could generate the observed empirical dynamics. A better understanding of the macroeconomic implications of dead-end duplications would shed light on the design of better macro industrial policies.

Finally, standard industrial policies, such as taxes, subsidies, and the degree of IP protection, would have nontrivial effects on the inefficiencies that we highlighted in this paper. For instance, the information-externality inefficiency could be alleviated by some R&D subsidy to the small firm, yet this could also exacerbate the dead-end replication problem. A detailed investigation of such policies would be an important area for future research.

Appendix: Proofs

A.1. Proof of Proposition 2

We begin with some useful observations. If the two firms start on the risky line together, continuing until a discovery is made, and then both switch to the safe line, their joint value is given by the following Bellman equation:

$$V = -\Lambda c dt + e^{-rdt}\left[\Lambda dt\left(\mu^0\Pi + w^{SS}\right) + (1-\Lambda dt)\,V\right],$$

which implies

$$V = \frac{\Lambda}{\Lambda + r}\left(\mu^0\Pi - c + w^{SS}\right). \tag{A.1}$$

This joint value can also be rewritten as

$$V = \frac{\Lambda\left(\mu^0\Pi - c\right)}{\Lambda + r} + \frac{\Lambda}{\Lambda + r}\frac{\Lambda\,(\pi - c)}{\Lambda + r} = w^{RR} + \frac{\Lambda}{\Lambda + r}w^{SS}. \tag{A.2}$$

Note that V consists of two parts. Firms first extract an expected payoff w^{RR} from the risky line, and meanwhile derive a flow payoff Λw^{SS} from the safe line with effective discounting $\Lambda + r$.

We relax the firms' decision problem by allowing reversibility; that is, they always have the option to restart a research line that they previously quit. This relaxed problem makes the computation of the continuation payoff easier. In the relaxed problem, the joint value $\hat{V}$ of the two firms can be derived from the following Bellman equation:

$$\hat{V} = \max \left\{ \begin{array}{l} \Lambda dt \left(\mu^0 \Pi + w^{SS}\right) e^{-rdt} - \Lambda c dt + (1 - \Lambda dt)\, \hat{V} e^{-rdt}, \\ \Lambda dt \left(\pi + w^{RR}\right) e^{-rdt} - \Lambda c dt + (1 - \Lambda dt)\, \hat{V} e^{-rdt}, \\ \lambda_1 dt \left(\mu^0 \Pi + w^{SS}\right) e^{-rdt} + \lambda_2 dt \left(\pi + w^{RR}\right) e^{-rdt} \\ \quad -\Lambda c dt + (1 - \Lambda dt)\, \hat{V} e^{-rdt}, \\ \lambda_2 dt \left(\mu^0 \Pi + w^{SS}\right) e^{-rdt} + \lambda_1 dt \left(\pi + w^{RR}\right) e^{-rdt} \\ \quad -\Lambda c dt + (1 - \Lambda dt)\, \hat{V} e^{-rdt}, \end{array} \right\}. \tag{A.3}$$

Here, the four terms on the RHS are the payoffs from strategies in which both firms start with the risky line, both firms start with the safe line, firm 1 starts with the risky line and firm 2 starts with the safe line, and firm 2 starts with the risky line and firm 1 starts with the safe line, respectively.

We claim that $\mu^0 \Pi + w^{SS} > \pi + w^{RR}$. This is because

$$\mu^0 \Pi + w^{SS} = \mu^0 \Pi + \frac{\Lambda}{\Lambda + r} (\pi - c) > \frac{\Lambda \mu^0 \Pi + r\pi + \Lambda \pi - \Lambda c}{\Lambda + r} = \pi + w^{RR}.$$

Note that the inequality follows from Assumption 1. Therefore, the first term on the RHS of (A.3) is the largest and hence $\hat{V} = \Lambda dt(\mu^0 \Pi + w^{SS}) e^{-rdt} - \Lambda c dt + (1 - \Lambda dt)\hat{V} e^{-rdt}$. This immediately implies that the optimal value of the relaxed problem, $\hat{V}$, is achieved by a strategy in which both firms start on the risky line. This strategy is feasible in the constrained problem where firms cannot switch back to a previously abandoned research line. Therefore, this strategy is optimal in the original problem, and the optimal value is given by equation (A.2),

$$V = w^{RR} + \frac{\Lambda}{\Lambda + r} w^{SS}.$$

This completes the proof.

A.2. Proof of Lemma 1

We conjecture that the differential equation has a solution of the following form:

$$\mu^t = \Psi(t) \equiv \frac{A}{1 + Bt}$$

where A and B are constants. Substituting the conjecture into equation (9) we obtain

$$\frac{-BA}{(1 + Bt)^2} = -\frac{A}{(1 + Bt)} \left(1 - \frac{A}{1 + Bt}\right) \frac{\lambda_{-n}}{1 + \lambda_{-n} t},$$

which reduces to $B + B\lambda_{-n}t = (1 - A)\lambda_{-n} + \lambda_{-n}Bt$. Equating the constant terms we obtain $B = (1 - A)\lambda_{-n}$. Moreover, we impose the boundary condition $\Psi(0) = \mu^0$. Then we obtain $A = \mu^0$ and $B = (1 - \mu^0)\lambda_{-n}$. This verifies our conjecture.

A.3. Proofs of Proposition 3 and Proposition 4

We proceed in four steps. In step 1, we characterize the stopping time T. In step 2, we show that both firms' stopping strategies are optimal. Last, step 3 proves the uniqueness.

Step 1: Characterization of the Stopping Time T. Suppose at time t, firm n's belief on the risky line is μ_n^t and its belief that its opponent, firm $-n$, is still on the risky line is β_n^t. Recall from equation (2) that w_n^{SS} is firm n's expected payoff from competing with firm $-n$ on the safe line,

$$w_n^{SS} = \frac{\lambda_n}{\Lambda + r}(\pi - c).$$

We define v_1^S as the value of firm 1 when it is alone on the safe line but anticipating that the strong firm 2 might switch to the safe line only after a discovery. Intuitively,

$$v_1^S = -\lambda_1 c dt + e^{-rdt}\left[\lambda_1 dt\pi + \lambda_2 dt w_1^{SS} + (1 - \Lambda dt)\, v_1^S\right],$$

which implies

$$v_1^S = \frac{\lambda_1(\pi - c) + \lambda_2 w_1^{SS}}{\Lambda + r} = w_1^{SS}\left(1 + \frac{\lambda_2}{\Lambda + r}\right).$$

In order for firm 1 to switch exactly at t, it must be that firm 1 is indifferent between switching at t or waiting until the next instant (we are assuming continuity of the value function and this will be true). The payoff from "stay on the risky research for another dt and then switch" is

$$\begin{aligned}&(1 - rdt)\,\lambda_1 dt\left\{\mu_1^t\left(\Pi + w_1^{SS}\right) + \left(1 - \mu_1^t\right)\left[b_1^{t+dt}v_1^S + \left(1 - b_1^{t+dt}\right)w_1^{SS}\right]\right\}\\&+ (1 - rdt)\,\beta_1^t\lambda_2 dt w_1^{SS} + (1 - rdt)\,(1 - \Lambda dt)\left[\beta_1^t v_1^S + \left(1 - \beta_1^t\right)w_1^{SS}\right] - \lambda_1 c dt.\end{aligned}$$

The first line is firm 1's discounted expected return when it makes a discovery on the risky line during $(t, t + dt)$. If the line is good, with probability μ_1^t, it leads to an immediate lump-sum payoff Π and a continuation payoff of competing in the safe research, w_1^{SS}; if the line is bad, the dead-end discovery gives rise to a 0 immediate payoff, but the expected continuation payoff depends on the position of the competitor. The second line is firm 1's discounted expected payoff in the case where the opponent firm 2 makes a discovery. It again depends on the position of firm 2. If firm 2 is on the risky line, which happens with probability β_1^t, firm 1 will compete with firm 2. If firm

2 is on the safe line, a discovery on the safe line indicates that the risky line is bad, and the game is over. The third line is firm 1's discounted expected payoff in the case of no discovery. The final line is the cost of researching.

The payoff from spending the next dt on the safe line and staying there forever is given by

$$(1-rdt)\,\lambda_1 dt\pi + (1-rdt)\,\beta_1^t\lambda_2 dt w_1^{SS} \\ + (1-rdt)\,(1-\Lambda dt)\left[\beta_1^t v_1^S + \left(1-\beta_1^t\right) w_1^{SS}\right] - \lambda_1 c dt.$$

The interpretation is similar to the previous case.

Therefore, by taking the limit, the indifference condition becomes

$$\mu_1^t\left(\Pi + w_1^{SS}\right) + \left(1-\mu_1^t\right)\left[b_1^t v_1^S + \left(1-b_1^t\right) w_1^{SS}\right] = \pi. \tag{A.4}$$

This condition carries the following intuition. At time t, spending an additional amount of time dt on either line delivers the same expected return conditional on an arrival of an outcome. To see this, note that the RHS is simply the expected return from the safe line. The left-hand side (LHS) is the expected return on the risky line. With probability μ_1^t, the line is good, in which case firm 1 receives the patent value Π and competes with firm 2 on the safe line and obtains w_1^{SS}. With the remaining probability $(1-\mu_1^t)$ the line is bad, in which case, firm 1 switches secretly to the safe line and obtains a payoff, depending on whether firm 2 is already on the safe line.

Therefore, the stopping time T is characterized by the following equation:

$$\mu_1^T\Pi + \left(1-\mu_1^T\right) b_1^T\left(v_1^S - w_1^{SS}\right) + w_1^{SS} = \pi. \tag{A.5}$$

From equations (8) and (10), we know that for $n = 1, 2$,

$$b_n^T = \frac{1}{1+\lambda_{-n}T} \quad \text{and} \quad \mu_n^T = \frac{\mu^0}{1+\left(1-\mu^0\right)\lambda_{-n}T}.$$

Hence

$$T = \underbrace{\frac{1}{\left(1-\mu^0\right)\lambda_2}}_{\text{Learning channel}}\underbrace{\left[\mu^0\left(\Pi + w_1^{SS}\right) + \left(1-\mu^0\right)v_1^S - \pi\right]}_{\text{Risky research premium}}\underbrace{\left(\frac{1}{\pi - w_1^{SS}}\right)}_{\text{Competition Channel}}$$

$$= \frac{1}{\left(1-\mu^0\right)\lambda_2}\left[\frac{\mu^0\Pi + \left(1-\mu^0\right)\left(v_1^S - w_1^{SS}\right)}{\left(\pi - w_1^{SS}\right)} - 1\right]$$

$$= \frac{1}{\left(1-\mu^0\right)\lambda_2}\left[\frac{\mu^0\Pi}{\pi}\frac{(r+\Lambda)}{\left(r+\Lambda-\lambda_1\frac{\pi-c}{\pi}\right)} - 1\right] + \frac{\lambda_1\frac{\pi-c}{\pi}}{(r+\Lambda)\left(r+\Lambda-\lambda_1\frac{\pi-c}{\pi}\right)}.$$

REMARK A.1 (Proposition 4). From the previous explicit expression for T, it is easy to check that T is increasing in μ^0 and Π, and decreasing in r, λ_2, and π. The comparative static relative to λ_1 is ambiguous.

Step 2: Best Responses of the Stopping Times in the Candidate Equilibrium. In this part, we show that the two firms' stopping times are best responses to each other.[23] In Step 4, after we have introduced the idea of an auxiliary problem, we shall show that the initial choices of the risky research line are mutual best responses in the candidate equilibrium.

Assume that firm 2 does not stop the risky research before a discovery. Recall that T is the unique solution of

$$\mu_1^t \Pi + \left(1-\mu_1^t\right) b_1^t \left(v_1^S - w_1^{SS}\right) + w_1^{SS} = \pi.$$

That is, T uniquely solves

$$\frac{\mu^0}{1+\left(1-\mu^0\right)\lambda_2 t}\Pi + \frac{\left(1-\mu^0\right)}{1+\left(1-\mu^0\right)\lambda_2 t}\left(v_1^S - w_1^{SS}\right) + w_1^{SS} = \pi.$$

We know that the LHS is monotonic decreasing in t. Hence if $t < T$, firm 1 strictly prefers to stay on the risky line, and if $t > T$, the firm strictly prefers to quit. Therefore, it is optimal for firm n to stop at $t = T$ before a discovery is made.

Now assume that firm 1 uses the stopping strategy characterized by T. Consider firm 2. There are two cases to consider.

Case 2.1: At $t \geq T$, firm 2's payoff conditional on being on the risky line in the candidate equilibrium is given by the recursion

$$\begin{aligned} V_2 = {} & -\lambda_2 c dt + (1 - r dt) \\ & \times \left[\lambda_2 dt \left(\mu_2^T \Pi + w_2^{SS}\right) + \lambda_1 dt \frac{\lambda_2}{r+\lambda_2}\left(\mu_2^T \Pi - c\right) + (1 - \Lambda dt) V_2\right]. \end{aligned}$$

Note that since $\mu_1^T \Pi - c \geq 0$ (otherwise, firm 1 would have already switched to the safe line before T), $\mu_2^T \Pi - c > 0$ by equation (10). Hence

$$V_2 = \frac{1}{r+\Lambda}\left[-\lambda_2 c + \lambda_2\left(\mu_2^T \Pi + w_2^{SS}\right) + \lambda_1 \frac{\lambda_2\left(\mu_2^T \Pi - c\right)}{r+\lambda_2}\right].$$

23. We assumed in our model that firms start on the risky research line. In the Online Appendix, we prove that in an alternative extensive form in which firms can choose which research lines to start with, the equilibrium characterized here continues to be an equilibrium, provided that certain parameter restrictions are made.

In order for firm 2 to stay on the risky research, we need $V_2 \geq w_2^{SS}$. Plugging in parameters, the sufficient condition can be simplified progressively as

$$
\begin{aligned}
-\lambda_2 c + \lambda_2 \left(\mu_2^T \Pi + w_2^{SS}\right) + \lambda_1 \frac{\lambda_2 \left(\mu_2^T \Pi - c\right)}{r + \lambda_2} &\geq (r + \Lambda)\, w_2^{SS} \\
\mu_2^T \Pi \left(1 + \frac{\lambda_1}{r + \lambda_2}\right) + w_2^{SS} - \frac{\lambda_1 c}{r + \lambda_2} &\geq \pi \qquad \text{(A.6)} \\
\mu_2^T \Pi - \pi + w_2^{SS} + \left(\mu_2^T \Pi - c\right) \frac{\lambda_1}{r + \lambda_2} &\geq 0.
\end{aligned}
$$

Note that at the time of the cutoff, the beliefs are such that $\mu_2^T > \mu_1^T$. A lower bound for μ_1^T is described as follows. Consider the same belief-updating procedure for firm 1, but now the payoffs are in such a way that the return on the risky line is higher and the return on the safe line is lower. This will give us a lower bound for μ_1^T since, in this environment, firm 1 will need a lower belief than the actual game to switch. To generate this payoff structure, assume that firm 1 does not face any competition on the risky line but faces competition with certainty on the safe line (continuing with the same belief updating). In that case the indifference condition in (A.5) reads as

$$\mu_1^{T^*} \Pi + w_1^{SS} = \pi$$

since $b_1^{T^*} = 0$. Therefore, we have

$$\mu_1^{T^*} = \frac{\pi - w_1^{SS}}{\Pi} < \mu_1^T < \mu_2^T.$$

Therefore, a sufficient condition for (A.6) is

$$\mu_1^{T^*} \Pi - \pi + w_2^{SS} + \left(\mu_1^{T^*} \Pi - c\right) \frac{\lambda_1}{r + \lambda_2} \geq 0.$$

Using the expression for $\mu_1^{T^*}$, the sufficient condition becomes

$$\frac{\left(\lambda_2 - \lambda_1\right)(\pi - c)}{r + \Lambda} + \frac{\lambda_1 (\pi - c)}{r + \Lambda} \geq 0.$$

This sufficient condition always holds.

Case 2.2: We need to show that firm 2 does not want to switch at any $t < T$. To this end, suppose, to the contrary, that firm 2 switches at $t < T$, while firm 1 follows the prescribed equilibrium strategy. Consider firm 2's response to the following strategy: firm 1 follows the candidate equilibrium strategy prescribed for firm 2.

If firm 2 has an incentive to switch at $t < T$ in the candidate equilibrium, it has an even stronger incentive to switch before t against the alternative strategy for firm 1 as prescribed previously. The reason is that the alternative strategy of firm 1 increases the competition on the risky line and reduces the competition on the safe line. We shall derive a contradiction as follows.

Given firm 1's alternative strategy, firm 2's belief goes down continuously over time before a discovery is observed, and hence there exists T_2 at which an indifference condition similar to (A.5) holds:

$$\pi = \mu_2^{T_2}\Pi + \left(1 - \mu_2^{T_2}\right) b_2^{T_2}\left(v_2^S - w_2^{SS}\right) + w_2^{SS}. \tag{A.7}$$

We claim that $T_2 > T$. To see this, suppose, to the contrary, that $T \geq T_2$. Then the following inequalities are immediate by definition: $\mu_2^{T_2} \geq \mu_2^T$, $\mu_2^{T_1} > \mu_1^T$, $(1-\mu_2^T)b_2^T > (1-\mu_1^T)b_1^T$, $v_2^S - w_2^{SS} > v_1^S - w_1^{SS}$, $w_2^{SS} > w_1^{SS}$. Using these inequalities, we derive from (A.7) that

$$\begin{aligned}\pi &= \mu_2^{T_2}\Pi + \left(1-\mu_2^{T_2}\right) b_2^{T_2}\left(v_2^S - w_2^{SS}\right)\\ &\quad + w_2^{SS} \geq \mu_2^T\Pi + \left(1-\mu_2^T\right) b_2^T\left(v_2^S - w_2^{SS}\right)\\ &\quad + w_2^{SS} > \mu_1^T\Pi + \left(1-\mu_1^T\right) b_1^T\left(v_1^S - w_1^{SS}\right) + w_1^{SS} = \pi.\end{aligned}$$

A contradiction.

Step 3: (Uniqueness) There are No Other Equilibrium Stopping Strategies when λ_2/λ_1 and $\mu^0\Pi/\pi$ are Large. Suppose to the contrary that there are other equilibria with stopping time T_1 and T_2. Since $\mu^0\Pi > \pi$, we know $T_1 > 0$ and $T_2 > 0$. We have two cases to consider.

Case 3.1: $T_2 < T_1 \leq \infty$.

We define $v_2^S(T_2, T_1)$ as the value of firm 2 at T_2 when it switches to the safe line but anticipating that firm 1 might switch to the safe line only after a discovery or at the random time τ_1.

First note that $T_2 < +\infty$ because of belief updating. In order for firm 2 to switch exactly at T_2, it must be that firm 2 is indifferent between switching at T_2 or waiting until the next instant and then switching. The payoff from "staying on the risky research line for another dt", is

$$\begin{aligned}(1-rdt)\,\lambda_2 dt &\left\{\mu_2^{T_2}\left(\Pi + w_2^{SS}\right) + \left(1-\mu_2^{T_2}\right)\right.\\ &\left.\times\left[b_2^{T_2+dt} v_2^S\left(T_2+dt, T_1\right) + \left(1-b_2^{T_2+dt}\right) w_2^{SS}\right]\right\}\\ &+ (1-rdt)\,\beta_2^{T_2}\lambda_1 dt w_2^{SS} + (1-rdt)(1-\Lambda dt)\\ &\times\left[\beta_2^{T_2+dt} v_2^S\left(T_2+dt, T_1\right) + \left(1-\beta_2^{T_2+dt}\right) w_2^{SS}\right] - \lambda_2 c dt.\end{aligned}$$

The payoff from "spend the next dt on the safe line and stay there forever," is given by

$$(1-rdt)\,\lambda_2 dt\pi + (1-rdt)\,\beta_2^{T_2}\lambda_1 dt w_2^{SS} + (1-rdt)\,(1-\Lambda dt) \\ \times \left[\beta_2^{T_2+dt} v_2^S\left(T_2+dt,T_1\right) + \left(1-\beta_2^{T_2+dt}\right) w_2^{SS}\right] - \lambda_2 c dt.$$

Therefore, by taking the limit, the indifference condition becomes

$$\mu_2^{T_2}\left(\Pi + w_2^{SS}\right) + \left(1-\mu_2^{T_2}\right)\left[b_2^{T_2} v_2^S\left(T_2,T_1\right) + \left(1-b_2^{T_2}\right) w_2^{SS}\right] = \pi,$$

or, equivalently,

$$\mu_2^{T_2}\Pi + \left(1-\mu_2^{T_2}\right) b_2^{T_2}\left[v_2^S\left(T_2,T_1\right) - w_2^{SS}\right] + w_2^{SS} = \pi. \tag{A.8}$$

Notice that $v_2^S(T_2,T_2) = w_2^{SS} \le v_2^S(T_2,T_1)$ for any $T_1 > T_2$. Then (A.8) gives us

$$\mu_2^{T_2}\Pi + w_2^{SS} \le \pi,$$

which is

$$T_2 \ge \frac{\mu^0\Pi - \left(\pi - w_2^{SS}\right)}{\left(\pi - w_2^{SS}\right)\left(1-\mu^0\right)\lambda_1}. \tag{A.9}$$

Now consider firm 1. Firm 1's belief on the risky line does not update after T_2, and its expected payoff is equivalent to that from staying on the risky line until a discovery, that is

$$\int_0^\infty e^{-(\Lambda+r)t}\left[\lambda_1\left(\mu_1^{T_2}\Pi - c + \frac{\lambda_1}{\Lambda+r}(\pi-c)\right) + \lambda_2\left(\frac{\lambda_1\left(\mu_1^{T_2}\Pi - c\right)}{\lambda_1+r}\right)\right]dt$$
$$= \frac{\lambda_1\left(\mu_1^{T_2}\Pi - c + \frac{\lambda_1}{\Lambda+r}(\pi-c)\right) + \lambda_2\left(\frac{\lambda_1\left(\mu_1^{T_2}\Pi - c\right)}{\lambda_1+r}\right)}{\Lambda+r}.$$

Since firm 1 has the option of competing on the safe line with firm 2, it must be that

$$\frac{\lambda_1\left(\mu_1^{T_2}\Pi - c + \frac{\lambda_1}{\Lambda+r}(\pi-c)\right) + \lambda_2\left(\frac{\lambda_1\left(\mu_1^{T_2}\Pi - c\right)}{\lambda_1+r}\right)}{\Lambda+r} \ge w_1^{SS} = \frac{\lambda_1(\pi-c)}{\Lambda+r}.$$

This condition can be simplified to

$$\mu_1^{T_2}\Pi - c \geq \frac{\lambda_1 + r}{\Lambda + r}\frac{\lambda_2 + r}{\Lambda + r}(\pi - c).$$

Hence,

$$T_2 \leq \frac{1}{\left(1-\mu^0\right)\lambda_2}\left[\frac{\mu^0\Pi}{\frac{\lambda_2+r}{\Lambda+r}\frac{\lambda_1+r}{\Lambda+r}(\pi - c) + c} - 1\right]. \tag{A.10}$$

Comparing (A.9) and (A.10), a contradiction will be derived if

$$\frac{\mu^0\Pi - \left(\pi - w_2^{SS}\right)}{\left(\pi - w_2^{SS}\right)\left(1-\mu^0\right)\lambda_1} > \frac{1}{\left(1-\mu^0\right)\lambda_2}\left[\frac{\mu^0\Pi}{\frac{\lambda_2+r}{\Lambda+r}\frac{\lambda_1+r}{\Lambda+r}(\pi - c) + c} - 1\right],$$

which is equivalent to

$$\mu^0\Pi\left[\frac{\lambda_2}{\frac{\lambda_1+r}{\Lambda+r}(\pi - c) + c} - \frac{\lambda_1}{\frac{\lambda_2+r}{\Lambda+r}\frac{\lambda_1+r}{\Lambda+r}(\pi - c) + c}\right] > \lambda_2 - \lambda_1. \tag{A.11}$$

First, since $\pi - c > 0$, we have

$$\begin{aligned}
&\frac{\lambda_2}{\frac{\lambda_1+r}{\Lambda+r}(\pi - c) + c} - \frac{\lambda_1}{\frac{\lambda_2+r}{\Lambda+r}\frac{\lambda_1+r}{\Lambda+r}(\pi - c) + c} \\
&\quad = \frac{\left(\lambda_2\frac{\lambda_2+r}{\Lambda+r} - \lambda_1\right)\frac{\lambda_1+r}{\Lambda+r}(\pi - c) + \left(\lambda_2 - \lambda_1\right)c}{\left[\frac{\lambda_1+r}{\Lambda+r}(\pi - c) + c\right]\left[\frac{\lambda_2+r}{\Lambda+r}\frac{\lambda_1+r}{\Lambda+r}(\pi - c) + c\right]} > \frac{\left(\lambda_2\frac{\lambda_2}{\Lambda} - \lambda_1\right)\frac{\lambda_1}{\Lambda}}{\pi}.
\end{aligned}$$

Hence a sufficient condition for (A.11) is

$$\frac{\mu^0\Pi}{\pi}\left[\left(\lambda_2\frac{\lambda_2}{\Lambda} - \lambda_1\right)\frac{\lambda_1}{\Lambda}\right] > \lambda_2 - \lambda_1.$$

This is guaranteed if

$$\frac{\lambda_2}{\lambda_1} > 2 \text{ and } \frac{\mu^0\Pi}{\pi} > \frac{\lambda_2 - \lambda_1}{\left(\lambda_2\frac{\lambda_2}{\Lambda} - \lambda_1\right)\frac{\lambda_1}{\Lambda}}.$$

Case 3.2: $T_1 \leq T_2 < \infty$. In this case, firm 2 does not update its belief after T_1 if it does not observe anything on the risky line. Therefore, for firm 2 to switch at $T_2 \geq T_1$, it must be that firm 2 is indifferent between switching at T_1 (competing with firm 1 on the safe line) and staying on the risky line (monopolizing the risky line with the option value of the safe line) at any $t \geq T_1$. Following the argument in the previous case, the indifference condition of firm 1 is

$$\mu_1^{T_1}\Pi + \left(1-\mu_1^{T_1}\right) b_1^{T_1}\left(v_1^S\left(T_1, T_2\right) - w_1^{SS}\right) + w_1^{SS} = \pi.$$

Recall that our equilibrium indifference condition is given by

$$\mu_1^T\Pi + \left(1-\mu_1^T\right) b_1^T\left(v_1^S - w_1^{SS}\right) + w_1^{SS} = \pi.$$

Since

$$b_n^T\left(1-\mu_n^T\right) = \frac{1-\mu^0}{1+\left(1-\mu^0\right)\lambda_{-n}T},$$

the LHS of the previous equation is strictly decreasing in T. Now suppose $T \le T_1$. Then it follows from $v_1^S > v_1^S(T_1, T_2)$ that

$$\begin{aligned}
\pi &= \mu_1^T\Pi + \left(1-\mu_1^T\right) b_1^T\left(v_1^S - w_1^{SS}\right) + w_1^{SS} \\
&\ge \mu_1^{T_1}\Pi + \left(1-\mu_1^{T_1}\right) b_1^{T_1}\left(v_1^S - w_1^{SS}\right) + w_1^{SS} \\
&> \mu_1^{T_1}\Pi + \left(1-\mu_1^{T_1}\right) b_1^{T_1}\left(v_1^S\left(T_1, T_2\right) - w_1^{SS}\right) + w_1^{SS} \\
&= \pi.
\end{aligned}$$

This is a contradiction. Hence $T > T_1$—that is, $\mu_2^T < \mu_2^{T_1}$.

In our equilibrium, firm 2 prefers to stay on the risky line after $T_1 > T$ upon no discovery and its belief is μ_2^T (since there is no updating between T and T_1). Hence

$$\frac{1}{\Lambda+r}\left[\lambda_2\left(\mu_2^T\Pi - c + w_2^{SS}\right) + \lambda_1\frac{\lambda_2}{\Lambda+r}\left(\mu_2^T\Pi - c\right)\right] \ge w_2^{SS}.$$

But at $t = T_1$ in the supposed equilibrium with stopping times $+\infty > T_2 > T_1$, we have for firm 2 (which is indifferent between staying on the risky line until a discovery or switching at T_1). Hence

$$\begin{aligned}
w_2^{SS} &= \frac{1}{\Lambda+r}\left[\lambda_2\left(\mu_2^{T_1}\Pi - c + w_2^{SS}\right) + \lambda_1\frac{\lambda_2}{\Lambda+r}\left(\mu_2^{T_1}\Pi - c\right)\right] \\
&> \frac{1}{\Lambda+r}\left[\lambda_2\left(\mu_2^T\Pi - c + w_2^{SS}\right) + \lambda_1\frac{\lambda_2}{\Lambda+r}\left(\mu_2^T\Pi - c\right)\right] \\
&= w_2^{SS},
\end{aligned}$$

where the strict inequality follows because $\mu_2^T < \mu_2^{T_1}$. This is a contradiction.

A.4. Proof of Lemma 2

No-Delay Condition. Suppose that firm n has an unreported dead-end discovery at time t (this discovery can be made right before t, or this discovery could have been made a while ago, which is off the equilibrium path). If firm n reveals the failure, then besides p_n^t it will get a continuation payoff

$$w_n^{SS} = \frac{\lambda_n}{\Lambda + r}(\pi - c).$$

Reporting immediately at t should lead to a higher payoff than delaying it to $t + h$ for any $h > 0$. That is,

$$\int_t^{t+h} e^{-(\Lambda+r)(\tau-t)}\left[-\lambda_n c + \lambda_n\left(\pi + p_n\right) + \lambda_{-n}\left(w_n^{ss} - p_{-n}\right)\right] d\tau \leq p_n + w_n^{SS} \tag{A.12}$$

holds for any $h > 0$. Since $p_n \geq 0$, the RHS of (A.12) is strictly positive. Therefore, whenever the integrand in the LHS is negative, then (A.12) holds trivially. If the integrand is strictly positive, the LHS is strictly increasing in h. Therefore, that (A.12) holds for any h is equivalent to

$$\left[-\lambda_n c + \lambda_n\left(\pi + p_n\right) + \lambda_{-n}\left(w_n^{ss} - p_{-n}\right)\right]\frac{1}{\Lambda + r} \leq p_n + w_n^{SS}.$$

If instead $-\lambda_n c + \lambda_n(\pi + p_n) + \lambda_{-n}(w_n^{ss} - p_{-n}) > 0$, then since the LHS of (A.12) is increasing in h, (A.12) is equivalent to

$$\left[-\lambda_n c + \lambda_n\left(\pi + p_n\right) + \lambda_{-n}\left(w_n^{ss} - p_{-n}\right)\right]\frac{1}{\Lambda + r} \leq p_n + w_n^{SS}.$$

This can be simplified into

$$\lambda_n\left(\pi - c - w_n^{SS}\right) \leq r\left(p_n + w_n^{SS}\right) + \lambda_{-n}\left(p_{-n} + p_n\right).$$

The intuition for this expression is as follows. By delaying, firm n loses the interest on $(p_n + w_n^{SS})$, and in the case of the opponent's discovery, firm n loses the transfer p_n and has to make an additional payment p_{-n} to the opponent. This is the RHS. Meanwhile, the firm makes an additional gain, which is equal to the benefit from monopolizing the safe line: $\lambda_n(\pi - c - w_n^{SS})$. Substituting w_n^{SS} into the previous expression and simplifying, we have

$$\frac{\lambda_{-n}\lambda_n}{\Lambda + r}(\pi - c) \leq \left(\lambda_{-n} + r\right)p_n + \lambda_{-n}p_{-n}. \tag{A.13}$$

No-Walk-Away upon Discovery of a Dead End. At any time, a firm should not leave the mechanism to start a decentralized competition. Let us denote firm n's value of walking away after the discovery of a failure at t as $v_{n,t}^S$, which is the value of monopolizing the safe line until firm $-n$ switches to the safe line. Note that for firm 1,

$v_{1,t}^S = v_{1,0}^S$ because firm 2 will never switch before a discovery. Therefore,

$$v_{1,0}^S = \int_0^\infty e^{-(\Lambda+r)t}\left[\lambda_1(\pi - c) + \lambda_2 w_1^{SS}\right]dt = w_1^{SS} + \frac{\lambda_2}{\Lambda + r}w_1^{SS}.$$

For firm 2, $v_{2,0}^S \geq v_{2,t}^S$ because firm 1 will switch at a finite time T even without a discovery. Therefore we can write $v_{2,0}^S$ as

$$\begin{aligned} v_{2,0}^S &= \int_0^T e^{-(\Lambda+r)t}\left[\lambda_2(\pi - c) + \lambda_1 w_2^{SS}\right]dt + \int_T^\infty e^{-(\Lambda+r)t}\lambda_2(\pi - c)\,dt \\ &= w_2^{SS} + \left[1 - e^{-(\Lambda+r)T}\right]\frac{\lambda_1}{\Lambda + r}w_2^{SS}. \end{aligned}$$

The value of sharing the information is $w_n^{SS} + p_n$. Therefore it must be that $w_n^{SS} + p_n \geq v_{n,0}^S$. Hence, we have another lower bound: $p_n \geq v_{n,0}^S - w_n^{SS}$. Therefore,

$$p_1 \geq \frac{\lambda_1\lambda_2}{(\Lambda + r)^2}(\pi - c) \text{ and } p_2 \geq \left[1 - e^{-(\Lambda+r)T}\right]\frac{\lambda_1\lambda_2}{(\Lambda + r)^2}(\pi - c). \quad \text{(A.14)}$$

Participation Constraint. The third condition is the participation constraint before any discovery. Let V_n^D be firm n's value in the decentralized market, $n = 1, 2$. Then the participation constraint is given by

$$V_n^D \leq \left\{\begin{array}{c} \mu^0 \int_0^\infty e^{-(\Lambda+r)t}\left[\lambda_n\left(\Pi - c + w_n^{SS}\right) + \lambda_{-n}w_n^{SS}\right]dt \\ + \left(1 - \mu^0\right)\int_0^\infty e^{-(\Lambda+r)t}\left[\lambda_n\left(p_n - c + w_n^{SS}\right) + \lambda_{-n}\left(w_n^{SS} - p_{-n}\right)\right]dt \end{array}\right\}.$$

The LHS is always V_n^D because when firm n walks away before any discovery, the game will resume as if the decentralized game has started at time $t = 0$ due to no updating until that point in the centralized market. This condition can be simplified to

$$\left(1 - \mu^0\right)\frac{\lambda_n p_n - \lambda_{-n}p_{-n}}{\Lambda + r} \geq V_n^D - \left[\frac{\lambda_n}{\Lambda + r}\left(\mu^0\Pi - c\right) + \frac{\Lambda}{\Lambda + r}w_n^{SS}\right].$$

By Proposition 2,

$$\frac{\lambda_n}{\Lambda + r}\left(\mu^0\Pi - c\right) + \frac{\Lambda}{\Lambda + r}w_n^{SS}$$

on the RHS is firm n's payoff V_n under full information sharing. Therefore, the condition can be rewritten as

$$\left(1 - \mu^0\right)\frac{\lambda_n p_n - \lambda_{-n}p_{-n}}{\Lambda + r} \geq V_n^D - V_n.$$

This expression is very intuitive. The LHS is the expected net transfer firm n receives from participating in the mechanism: there will be transfer only when the risky line has a dead end that occurs with a prior probability $(1-\mu^0)$; on the equilibrium path, the belief will never update because of full information sharing; firm n receives a transfer p_n at a rate λ_n and makes a transfer p_{-n} at a rate λ_{-n}, and hence the discounted value of the net transfer on a dead-end line is $(\lambda_n p_n - \lambda_{-n} p_{-n})/(\Lambda + r)$. The RHS is the value firm n gives up by participating in the mechanism: it obtains a value V_n under full information sharing enforced by the mechanism, but V_n^D in a decentralized market:

$$\lambda_n p_n - \lambda_{-n} p_{-n} \geq \frac{\Lambda + r}{1-\mu^0}\left(V_n^D - V_n\right).$$

This condition holds for $n = 1, 2$, and hence, we obtain an upper bound and a lower bound for $\lambda_1 p_1 - \lambda_2 p_2$:

$$\underline{K} \leq \lambda_1 p_1 - \lambda_2 p_2 \leq \bar{K}.$$

$$\text{where} \quad \underline{K} \equiv \frac{\Lambda + r}{1-\mu^0}\left(V_1^D - V_1\right) \quad \text{and} \quad \bar{K} \equiv \frac{\Lambda + r}{1-\mu^0}\left(V_2 - V_2^D\right).$$

It is feasible only when $\underline{K} \leq \bar{K}$. This condition is equivalent to $V_1^D + V_2^D \leq V_1 + V_2$. The RHS is the first-best joint payoff under full information. The LHS is the sum of values of the firms in the decentralized economy. Clearly, this condition is always satisfied.

References

Acemoglu, Daron (2008). *Introduction to Modern Economic Growth*. Princeton University Press.

Acemoglu, Daron and Ufuk Akcigit (2012). "Intellectual Property Rights Policy, Competition and Innovation." *Journal of the European Economic Association*, 10, 1–42.

Aghion, Philippe, Christopher Harris, Peter Howitt, and John Vickers (2001). "Competition, Imitation and Growth with Step-by-step Innovation." *Review of Economic Studies*, 68, 467–492.

Aghion, Philippe and Peter Howitt (1998). *Endogenous Growth Theory*. MIT Press.

Arrow, Kenneth (1962). "Economic Welfare and The Allocation of Resources for Invention." In *The Rate and Direction of Inventive Activity: Economic and Social Factors*, edited by R. R. Nelson. NBER, pp. 609–626.

Bergemann, Dirk and Ulrich Hege (1998). "Venture Capital Financing, Moral Hazard, and Learning." *Journal of Banking and Finance*, 22, 703–735.

Bergemann, Dirk and Ulrich Hege (2005). "The Financing of Innovation: Learning and Stopping." *RAND Journal of Economics*, 36, 719–752.

Bergemann, Dirk and Juuso Välimäki (1996). "Learning and Strategic Pricing." *Econometrica: Journal of the Econometric Society*, 64, 1125–1149.

Bergemann, Dirk and Juuso Välimäki (2000). "Experimentation in Markets." *Review of Economic Studies*, 67, 213–234.

Bhattacharya, Sudipto, Kalyan Chatterjee, and Larry Samuelson (1986). "Sequential Research and the Adoption of Innovations." *Oxford Economic Papers*, 38, 219–243.

Bhattacharya, Sudipto, Jacob Glazer, and David EM Sappington (1992). "Licensing and the Sharing of Knowledge in Research Joint Ventures." *Journal of Economic Theory*, 56, 43–69.

Bhattacharya, Sudipto and Sergei Guriev (2006). "Patents vs. Trade Secrets: Knowledge Licensing and Spillover." *Journal of the European Economic Association*, 4, 1112–1147.

Blundell, Richard, Rachel Griffith, and John Van Reenen (1999). "Market Share, Market Value and Innovation in a Panel of British Manufacturing Firms." *Review of Economic Studies*, 66, 529–554.

Bolton, Patrick and Christopher Harris (1999). "Strategic Experimentation." *Econometrica*, 67, 349–374.

Bonatti, Alessandro and Johannes Hörner (2011). "Collaborating." *American Economic Review*, 101, 632–663.

Chatterjee, Kalyan and Robert Evans (2004). "Rivals' Search for Buried Treasure: Competition and Duplication in R&D." *RAND Journal of Economics*, 15, 160–183.

Das, Kaustav (2014). "Strategic Experimentation with Competition and Private Arrival of Information." Technical report.

d'Aspremont, Claude, Sudipto Bhattacharya, and Louis-André Gérard-Varet (2000). "Bargaining and Sharing Innovative Knowledge." *Review of Economic Studies*, 67, 255–271.

Fershtman, Chaim and Ariel Rubinstein (1997). "A Simple Model of Equilibrium in Search Procedures." *Journal of Economic Theory*, 72, 432–441.

Fudenberg, Drew and Jean Tirole (1985). "Preemption and Rent Equalization in the Adoption of New Technology." *Review of Economic Studies*, 52, 383–401.

Gill, David (2008). "Strategic Disclosure of Intermediate Research Results." *Journal of Economics and Management Strategy*, 17, 733–758.

Green, Jerry R and Suzanne Scotchmer (1995). "On the Division of Profit in Sequential Innovation." *RAND Journal of Economics*, 20–33.

Grossman, Gene M and Elhanan Helpman (1993). *Innovation and Growth in the Global Economy*, MIT Press.

Halac, Marina, Navin Kartik, and Qingmin Liu (2013). "Optimal Contracts for Experimentation." Working Paper, Department of Economics, Columbia University.

Halac, Marina, Navin Kartik, and Qingmin Liu (2014). "Contests for experimentation." Working Paper, Department of Economics, Columbia University.

Heidhues, Paul, Sven Rady, and Philipp Strack (2014). "Strategic Experimentation with Private Payoffs." University of Bonn Working Paper.

Hopenhayn, Hugo A and Francesco Squintani (2011). "Preemption Games with Private Information." *Review of Economic Studies*, 78, 667–692.

Hörner, Johannes and Larry Samuelson (2013). "Incentives for Experimenting Agents." *RAND Journal of Economics*, 44, 632–663.

Jensen, Richard (1981). "Adoption and Diffusion of an Innovation of Uncertain Probability." *Journal of Economic Theory*, 27, 182–193.

Jovanovic, Boyan and Yaw Nyarko (1996). "Learning by Doing and the Choice of Technology." *Econometrica*, 64, 1299–1310.

Jovanovic, Boyan and Rafael Rob (1990). "Long Waves and Short Waves: Growth through Intensive and Extensive Search." *Econometrica*, 58, 1391–1409.

Katz, Michael L and Carl Shapiro (1987). "R and D Rivalry with Licensing or Imitation." *American Economic Review*, 77(3), 402–420.

Keller, Godfrey, Sven Rady, and Martin Cripps (2005). "Strategic Experimentation with Exponential Bandits." *Econometrica*, 73, 39–68.

Lerner, Joshua (2012). *The Architecture of Innovation: The Economics of Creative Organizations*. Harvard Business Press.

Manso, Gustavo (2011). "Motivating Innovation." *Journal of Finance*, 66, 1823–1860.

Moscarini, Giuseppe and Francesco Squintani (2010). "Competitive Experimentation with Private Information: The Survivor's Curse." *Journal of Economic Theory*, 145, 639–660.

Murto, Pauli and Juuso Välimäki (2011). "Learning and Information Aggregation in an Exit Game." *Review of Economic Studies*, 78, 1426–1461.

Nanda, Ramana and Matthew Rhodes-Kropf (2012). "Innovation and the Financial Guillotine." Harvard Business School Entrepreneurial Management Working Paper 13-038.

Nickell, Stephen J (1996). "Competition and Corporate Performance." *Journal of Political Economy*, 104, 724.

Nielsen, Michael (2012). *Reinventing Discovery: The New Era of Networked Science*. Princeton University Press.

Pastorino, Elena (2011). “Learning in Labor Markets and Job Mobility.” Working paper, University of Minnesota.

P*h*RMA (2011). *Pharmaceutical Industry: Profile 2011*, The Pharmaceutical Research and Manufacturers of America, Washington.

Scotchmer, Suzanne (2004). *Innovation and Incentives*, MIT Press.

Singer, Natasha (2009). “Seeking A Shorter Path to New Drugs.” *New York Times*. Accessed online at: http://www.nytimes.com, November 14, 2009.

Singh, Simon (1997). *Fermat’s Enigma: The Epic Quest to Solve the World’s Greatest Mathematical Problem*, Anchor Books.

Spiegel, Yossi (2008). “Licensing Interim R&D Knowledge.” Tel Aviv University Working Paper.

The White House, Office of the Press Secretary (2000). *Joint Statement by President Clinton and Prime Minister Tony Blair of the UK*. Accessed online at: http://clinton4.nara.gov/WH/EOP/OSTP/html/00314.html, March 14, 2000.

Thomas, Caroline (2011). “Experimentation with Congestion.” Working paper, University of Texas at Austin.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher’s web site:

Online Appendix