

NBER WORKING PAPER SERIES

THE ROLE OF INFORMATION IN COMPETITIVE EXPERIMENTATION

Ufuk Akcigit
Qingmin Liu

Working Paper 17602
<http://www.nber.org/papers/w17602>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2011

The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research. We are grateful to Daron Acemoglu and Andrzej Skrzypacz for their invaluable comments and suggestions. We also thank Alessandro Bonatti, Kalyan Chatterjee, Benjamin Golub, Christopher Harris, Hugo Hopenhayn, Johannes Horner, Matthew Jackson, Nicolas Klein, Dirk Krueger, Antonio Merlo, Matthew Mitchell, Andrew Postlewaite, Joel Sobel, Francesco Squintani, and the participants at the AEA 2011 winter meeting and the 2011 NSF/CEME Decentralization Conference for helpful discussions. We thank Francesc Dilme and Zehao Hu for excellent research assistance. The usual disclaimer applies. Akcigit gratefully acknowledges the financial support of NBER Innovation Policy Research Grant.

The online appendix can be found at <http://www.sas.upenn.edu/~uakcigit/web/Research.html>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2011 by Ufuk Akcigit and Qingmin Liu. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Role of Information in Competitive Experimentation
Ufuk Akcigit and Qingmin Liu
NBER Working Paper No. 17602
November 2011, Revised November 2011
JEL No. D83,D92,O31

ABSTRACT

Technological progress is typically a result of trial-and-error research by competing firms. 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 dead ends. Consequently, technological progress is slowed down, and the society benefits from innovations with delay, if ever. To study this prevalent problem, we build a tractable two-arm bandit model with two competing firms. The risky arm could potentially lead to a dead end and the safe arm introduces further competition to make firms keep their dead-end findings private. We characterize the equilibrium in this decentralized environment and show that the equilibrium necessarily entails significant efficiency losses due to wasteful dead-end replication and a flight to safety – an early abandonment of the risky project. Finally, we design a dynamic mechanism where firms are incentivized to disclose their actions and share their private information in a timely manner. This mechanism restores efficiency and suggests a direction for welfare improvement.

Ufuk Akcigit
Department of Economics
University of Pennsylvania
3718 Locust Walk, #445
Philadelphia, PA 19104
and NBER
uakcigit@econ.upenn.edu

Qingmin Liu
Columbia University
1022 International Affairs Building
Mail Code 3308
420 West 118th Street
New York, NY 10027
qingmin@econ.upenn.edu

Contents

1	Introduction	1
2	Related Literature	4
3	Model	7
3.1	Basic Environment	8
3.2	The Safe Arm	10
3.2.1	Monopoly	11
3.2.2	Cooperation	11
3.2.3	Competition	12
4	Equilibrium Analysis of the Model	12
4.1	Monopoly	13
4.2	Cooperation: Planner's Problem	13
4.3	Competition in a Decentralized Market	14
4.3.1	Learning and Private Beliefs	15
4.3.2	Equilibrium	18
5	Numerical Example	20
5.1	Summary Statistics	20
5.2	Types of Inefficiencies: Dead End and Early Switching	24
6	Mechanism Design for Efficient Competition	25
6.1	Feasible Mechanisms	26
6.2	Incentives	27
6.2.1	No-delay Condition	27
6.2.2	No Walk-away upon Discovery of a Dead End	28
6.2.3	Participation Constraint	29
6.3	Efficient Mechanism	30
6.4	Minimum Implementable Transfers	31
7	Concluding Discussion and Extensions	32
7.1	State-dependent Arrival Rate	32
7.2	Strategic Patenting	33
7.3	Macroeconomic Applications	34
8	References	36
A	Proof of Proposition 2	39
B	Proof of Lemma 1	40
C	Proofs of Theorem 1 and Proposition 3	41
D	Details of the Numerical Example	49

1 Introduction

Innovation is a risky process in which the exact path to success is unknown. Therefore, many potential innovators go through trial-and-error experimentation that leads to high R&D costs. Pharmaceutical companies offer a typical case in point. According to a report by Pharmaceutical Research and Manufacturers of America (PhRMA, 2011), developing a drug can cost more than \$1 billion and take 10 to 15 years. The final cost of a drug arises mostly from early failed attempts to develop it. All firms in competition with each other to develop a particular drug typically follow similar paths: they try out and then give up on similar compounds due to toxicity or inefficacy. Yet, firms do not reveal to each other their research activities. In particular, they do not share information about which exploratory paths they have pursued and have been proven to be fruitless. As a result many firms waste years and millions on projects that their competitors have already found to be dead ends.

This problem is common and severe in many industries where research progresses mostly through trial-and-error.¹ While an extensive literature on R&D and innovation focuses on the incentives for and the impacts of *successful* innovations, this paper turns to the dark side of the picture and studies the *failed* innovation attempts that incur huge costs, both for firms and for society overall. Our goal is to analyze the private and social values of dead-end research paths and to understand firms' incentives to keep them private, that is, to not reveal dead-end paths to their competitors. In addition, we want to discuss a mechanism that can potentially undo this inefficiency.

Duplicative R&D efforts have attracted attention both in academic and industrial spheres (see, Kortum, 1993; Tufts Report, 2009). Delving deeper into the details of the modern drug research process would help us understand the problem of duplication better. PhRMA (2011) describes the process in great detail. The research for a drug typically starts with the scientific diagnosis of the proteins causing a disease. Often, the aim of the sought-after drug is to inhibit some protein activity causing the physiological harm or to stimulate some protein activity that is missing. The next step is to discover a chemical that will bind to the *target* protein either to inhibit it or to help it function as it normally should. This is where the *trial-and-error* procedure starts. Companies try out about 5,000 to 10,000 chemical compounds (drug candidates) to see if they bind to the target protein. Once the promising drug candidates are identified, *preclinical testing* starts. Out of those thousands of candidates, only about 250 make it to the

¹We will nevertheless mostly refer to pharmaceutical research, to simplify the exposition.

preclinical stage, while the rest are simply recorded in the company's private database. The successful subset of candidates is tested on animals (as well as pregnant animals to test the effect on pregnancy²), for toxicity and efficacy. Out of those 250 chemicals, approximately 5 successfully move on to the *clinical trials* stage, in which they are tested on human beings. For the ones that pass those clinical trials successfully, the company files a New Drug Approval application with the Food and Drug Administration. This whole process takes about 10 to 15 years and can cost more than \$1 billion, most of which are clearly spent on the trial-and-error efforts, according to PhRMA (2011). These are only the accounting costs on the firms' balance sheets. The economy endures further cost such as delayed cures for the patients and a slowdown in the growth of the entire sector.

Two key features distinguish the above R&D process from the ones considered by the literature to date. First, if we think about each of the initial drug candidates as a research line in itself, it becomes clear that this line could lead to a good outcome or to a dead end. The existence of a positive reward of a particular research line is highly uncertain. Second, a firm's research activities are confidential and the dead-end outcomes are kept private within the firm because of competitive pressures, even though society would benefit from their revelation. Publicizing a dead-end outcome makes the research line obsolete for everyone, while disclosing private actions provides valuable *indirect inference* to the opponents. These two features make this type of R&D competition unique and different from the previous R&D models in the literature in which typically the existence of a positive outcome is certain, although its arrival rate is stochastic. These two aforementioned new features will be the building blocks of our analysis, and we will ask the following related questions: What are the implications of the two features on firms' innovation strategies? What is the cost of sharing information to a company that has discovered that a particular project is a dead end? What are the potential inefficiencies in an R&D competition setup in which firms do not disclose their failed attempts? What could be a mechanism to improve efficiency? These questions are central to policy debates on intellectual property rights and R&D. Part of our agenda is to construct a tractable model that could lay the micro foundation for an endogenous growth model with asymmetric information.

To study the aforementioned questions, we build a parsimonious two-arm bandit

²This entails a non-material cost that is difficult to measure in dollars. These tests on pregnant animals including monkeys have been opposed by many groups, and scientists have been trying to develop alternative, potentially more costly (in monetary terms), testing methodologies.

model with two asymmetric firms that differ in their arrival rates of innovation. Each firm can research at most one research line at a time and has to pay a cost $c > 0$ per arrival rate and per unit of time on the research. The arrival of outcomes in both lines follow Poisson arrival processes. Though the lines share the stochastic nature of outcome arrivals, they differ in one crucial aspect. The safe research is commonly known to deliver a one-time lump-sum payoff $\pi > 0$ upon an outcome arrival. The risky research is ex-ante more profitable, yet has an additional uncertainty. In particular, 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 (e.g., market value of a drug), while a bad outcome reveals that the risky research line is a dead-end, in which case the payoff is simply 0. Certain approaches to the cures of HIV or various cancers are potential examples for the risky research line, and research on incremental improvements on existing drugs are examples of the safe research line. The lump-sum payoff associated with the good outcome arises from the *publicly observable* patent for that particular drug and the resulting monopoly power for that market. We ignore consumers' payoffs in this analysis of firm competition, the inclusion of which would only strengthen our results on efficiency loss.

Firms share a common prior on the probability of the risky research arm being good. Both firms start on this risky research line. If the arriving outcome is good, the firm obtains a publicly observable payoff Π . However, if the outcome is a dead end, the firm quits this research arm and switches to safe research. The key feature in this case is that the rival cannot observe the firm's switching action. Therefore firms form belief both on the nature of the risky arm and the rival's position. The rival may continue with its own research without knowing that the arm is a dead end. This is the first inefficiency that emerges in our model. We call this the *dead-end inefficiency*. The second inefficiency arises due to equilibrium belief updating. At any point in time, three events can occur in a firm: the firm (*i*) receives a good outcome and patents it, (*ii*) receives a bad outcome and secretly quits, (*iii*) does not receive any outcome. Since only the first case is observable to a competitor, when a firm does not observe any outcome from its rival, it will update downwards its belief about the success of the research arm. As a result, a firm could eventually quit the research line and switch to the safe research, even though it has neither itself discovered any outcome, nor observed any patent from its rival. This will be a second channel of inefficiency, in the case where the research arm is actually a good arm. We call this the *early-switching inefficiency*.

Note that our framework features both private actions and private outcomes. Hence the solution of the model requires keeping track of two payoff-relevant beliefs: one about

the nature of the risky research and another about the position of the competitor. We achieve the tractability in this environment by focusing on a pure strategy equilibrium. We characterize a pure strategy equilibrium and show that it is unique if the game features enough asymmetry across the firms and the research lines. We identify the aforementioned inefficiencies in this environment and find that the asymmetric firms generate different inefficiencies. While both firms endure dead-end inefficiency, only the weak firm (the firm with a lower arrival rate) creates early-switching inefficiency. Our model suggests that a seemingly negligible amount of competition on the safe arm generates a drastic welfare loss. Next, we propose a dynamic mechanism that could undo these inefficiencies. This mechanism involves a third party that collects monetary installments, ex-ante, and rewards the report of failed attempts as time progresses in an incentive-compatible way. As a result, firms are incentivized to participate in the mechanism at any point in time, share their dead-end findings *without any delay* upon their discovery, and follow the first-best decision rules. The basic idea we try to convey is that we should also consider *rewarding the failed attempts* in order to improve efficiency. Our paper thus complements and contrasts with the literature on *patents for successful innovations* as a reward mechanism. Notice that while private industries currently reward only profitable, positive outcomes, “patents for dead-end discoveries” already exist in many academic professions that publish the impossibility results.³

We view our main objectives in this paper as follows. First, the economic problem we consider here is a general one that applies to many industries and has significant welfare implications. We hope that our paper would draw attention to this practical and fundamental problem and would promote further investigations on implementable institutional design to remedy this problem. Second, our tractable model can serve as a workhorse for further investigations and can be enriched to consider alternative environments. We offer more details on this in the conclusion. In what follows, we will place our paper in the literature and elaborate our contributions in more details.

2 Related Literature

Our paper contributes to the broad literature on R&D races. The type of duplication that arises in that literature is very different from the duplication we capture here. For instance, Harris and Vickers (1985, 1987); Aghion et al. (2001); Acemoglu and Akcigit (2011) consider an R&D competition model in which the technology gap between

³We thank Matthew Jackson for pointing this out to us.

the competing firms is endogenously determined through the R&D investments of the leader and follower. While the technology leader's successful R&D pushes forward the technology frontier, the follower's successful R&D effort replicates the steps that were previously already taken by the leader. As a result, the follower's R&D effort is spent on wasteful duplications. In our model, competing firms replicate each other's dead-end results as opposed to the successful findings, and the R&D race is efficient as long as private information is made public.

There are related R&D race models with social learning. Chatterjee and Evans (2004) offer a fully dynamic two-arm bandit model of R&D rivalry.⁴ 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. Moreover, they assume that both actions and outcomes are perfectly observable. As they stated, the central trade-off is that an agent wants to take a different arm from his opponent to reduce the possibility of simultaneous discovery (which leads to a low payoff due to Bertrand competition), while doing so increases the chance of leaving the opponent exploiting the correct arm. In contrast, the central trade-off in our model is very different. Simultaneous discovery is not an issue in our continuous-time context, and it is precisely the private observations and private strategies in conjunction with competition that drive our results. Fershtman and Rubinstein (1997) investigate a two-stage model 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.⁵ In this model, the central theme is that an agent wants to preempt his opponent by opening a box before his opponent. Since players face the same set of boxes and the search order is chosen once and for all, an equilibrium must involve randomization over the orders of boxes. As a result of randomization, the most likely box is not searched first. There is indeed 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. Moscarini and Squintani (2010) consider a one-arm bandit competitive experimentation model with publicly observable stopping decisions and initial private information. There is no information arrival and hence no dead-end discovery either,

⁴See also Bhattacharya, Chatterjee and Samuelson (1986) for a model of dynamic adoption of innovation with Gaussian signals where firms observe each other's actions. This paper thus precedes the bandit literature that deals with observable actions and private signals which we shall discuss shortly.

⁵In their model, agents could also choose to group several boxes together and the group size is also a choice variable.

but interesting social learning takes place from observing the opponent staying in the game, and a quitting decision of the opponent also reveals his initial private signals.⁶ Both our motivation and our model differ from theirs. In our two-arm bandit model, actions and the (bad) outcome are private, and it is the competitive incentive on the safe arm and new information arrival that affect the experimentation and learning on the risky arm.

In this paper, we study the social value of failed attempts and suggest a dynamic mechanism to also *reward the failed attempts* in order to prevent wasteful duplication of dead-end projects. Our paper thus complements the literature on patents as a reward mechanism to *successful innovations*. The main purpose of a patent is to provide ex-post monopoly power so that agents can engage in costly innovation efforts ex-ante (e.g., Arrow (1962); Reinganum (1982); Scotchmer (1991); Aghion and Howitt (1992)). The main focus in this literature has been the trade-off between the length and width of the patent protection. Acemoglu, Bimpikis, and Ozdaglar (2011) consider a model in which firms receive private signals on the success probability of research projects and decide which one to implement. They show that patents can prevent inefficient delays. In an R&D race model, Acemoglu and Akcigit (2011) numerically show that the design of intellectual property rights can be used to provide additional incentives by providing stronger protections to the more advanced innovators. Hopenhayn and Mitchell (2011) take a mechanism design approach and completely characterize the optimal mechanism in a model with recurrent innovators. See also Kremer (1998), Hopenhayn, Llobet and Mitchell (2006) and Hopenhayn and Mitchell (2011) for related investigations.

Finally, our paper contributes to the strategic bandit literature.⁷ Multi-agent experimentation in teams has been studied in Bolton and Harris (1999); Keller, Rady, and Cripps (2005); Strulovici (2010), and Bonatti and Hörner (2011). Klein and Rady (2011) also consider a negatively correlated two-arm bandit problem. Free-riding is a common feature in these models. This leads to an inefficient level of experimentation and often mixed strategy equilibria. In contrast, in our model, we focus on applications with a winner-takes-all competition. The only observable event is one's rival's success, which lowers one's payoff. Therefore the free-riding motive does not emerge in our model. Our model also differs from Thomas (2011). In her model, two players share a safe option that can be taken by at most one player at a time, and each player

⁶Murto and Välimäki (2011) study a related social learning model with common payoffs.

⁷See Bergemann and Välimäki (2008) for a survey. See also Moscarini and Smith (2001) for their treatment of a single decision maker bandit problem and its connection with R&D models.

has an independent risky option. Interestingly, the congestion on the safe arm leads to inefficient experimentation on the risky option. Similar to the previous studies, there is no dead-end information in her model, and both actions and outcomes are publicly observable. Rosenberg, Solan, and Vieille (2007); Hopenhayn and Squintani (2011); and Murto and Välimäki (2011) consider bandit problems in which outcomes are not observable but actions are. In our paper, we consider a two-arm bandit problem in which the actions are unobservable, and outcomes are only partially observable. As a result, firms have to form beliefs both about their rivals' actions and about the existence of a positive risky payoff. The inefficient level of learning and experimentation is not due to free-riding; rather it is due to uncertainty about the type of an opponent's discovery, if any. We need to emphasize that in free-riding bandit papers, early switching is due to the assumption that no news is bad news, while in our model, it arises *endogenously* through competition – in fact, our model does not generate early switching under perfect information. Another contrast to most of the two-arm bandit models is that the safe arm in our model features stochastic arrival; this plays an important role as the arrival could reveal information about a firm's past private observation on the risky arm, which makes the two arms correlated through this information channel.

More importantly, the aforementioned strategic bandit models study fixed games with specific assumptions on observability of actions and outcomes. Our paper is the first one to take a mechanism design approach to study efficient information sharing in a bandit environment.

The rest of the paper is structured as follows. Section 3 outlines the model. Section 4 characterizes the equilibrium in a decentralized market. Section 5 provides a numerical example. Section 6 studies the mechanism design for information sharing. Section 7 concludes and also provides a discussion of potential extensions.

3 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 *both* arms. In the sequel, we attempt to offer the simplest possible model that captures the essence of the central trade-offs in such market environments.

3.1 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 firms start the game with a competition on the risky arm.⁸ The arrival of outcomes in both lines follow Poisson arrival processes. The safe research is commonly known to deliver a one-time lump-sum payoff $\pi > 0$ upon an outcome arrival. 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.⁹

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 its actions, but it cannot return to the research line it had left. This irreversibility assumption simplifies the analysis of inference/belief-updating without affecting our main focus; it comes at a cost: the calculation of continuation payoff is more involved.¹⁰

⁸In the online appendix, we extend the model by allowing firms to choose simultaneously at $t = 0$ which arm to start to with, and in particular, they could start with the safe arm and switch to the risky arm later. This extension introduces interesting strategic considerations, though they are not directly related to our motivation.

⁹The 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 a realistic condition also from an empirical point of view.

¹⁰Reversibility usually does not play any role in stopping games such as in Bonatti and Hörner (2011). This is not the case in our model, because the players in our game have three options; players' payoffs depend on the arms they choose and the stochastic arrival reveals information. Different versions of irreversibility assumptions appear elsewhere in the literature such as Rosenberg, Solan, and Vieille

The firm’s research activity is private and unobservable to the public. However, a successful discovery is public.¹¹ Therefore, a firm is uncertain of 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.

For the purpose of formal analysis, we endow the continuous-time game with two private stages $k = 0, 1$ for each firm. This is an adoption of the “public stage” idea proposed by Murto and Välimäki (2011) into our setup with unobservable actions to overcome the well-known modelling issue in continuous time.¹²

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 to firm n . New information takes one of the following three forms:

1. firm n makes a discovery on the risky research arm,
2. firm n observes a good-outcome discovery from its competitor on the risky arm, and
3. firm n observes a discovery from its competitor on the safe arm.

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.

(2007); Murto and Välimäki (2011); and Acemoglu, Bimpikis, and Ozdaglar (2011).

¹¹For 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 7.

¹²We allow a firm to react immediately, without a lag, to new information it obtains either by making discovery on its own, or observing potential good discoveries by its opponent. This creates a well-known modelling issue of timing of events in continuous time. The standard approach adopted in the exponential bandit literature is to focus on Markov strategies that depend on the beliefs over the risky arm, 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 not only depends 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-arm problem with irreversibility, we need to keep track of the arms that have been visited in the past (this is not necessary in a one-arm problem, as switching arms ends the game).

If firm n enters its private stage $k = 1$ at $t = T_{n,0}$ when its stopping time expires without observing new information arrival, 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 a new information arrival, 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 arm if it has not abandoned it yet; while in the former case, firm n voluntarily gives up the risky arm 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 information arrives. Note that the game only consists of at most two private stages for each firm because an observable discovery will remove a research line from the choice set. We focus on perfect Bayesian equilibrium in pure strategies.¹³

Below, we summarize the notation that appears frequently in the main text to facilitate the reading of the paper.¹⁴

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

TABLE 1

3.2 The Safe Arm

The core of our idea is that competition in the safe research prevents the disclosure of socially efficient information regarding the risky research line. To understand the

¹³In contrast, pure strategy equilibria usually do not exist in bandit problems with free-riding.

¹⁴In choosing this notation, the superscript SS indicates there are *two* firms on the *safe* arm; the superscript S indicates only *one* firm is on the *safe* arm. The subscript n indicates that the profit is attributed to firm n .

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 findings here will be used later to determine the equilibrium continuation payoffs. In the sequel, we characterize the strategic behaviour in three different market structures: monopoly, cooperation, and competition.

3.2.1 Monopoly

Write firm n 's monopolistic value from the safe arm as w_n^S . Assuming the firm's strategy is to work on the arm until a discovery is made, we can express w_n^S recursively using the following continuous time Hamilton-Jacobi-Bellman (HJB) equation:

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

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

The HJB equation immediately gives us

$$w_n^S = \frac{\lambda_n}{\lambda_n + r} (\pi - c). \quad (2)$$

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 arm if $\pi > c$.

Assumption 2 $\pi > c$.

It also transpires from the monotonicity of $\frac{\lambda_n}{\lambda_n + r}$ in λ_n that the strong firm enjoys larger monopolistic profits.

3.2.2 Cooperation

Next, we consider the cooperative scheme in which firms maximize their joint value, w^{SS} . The HJB equation is

$$w^{SS} = -\Lambda c dt + e^{-rdt} [\Lambda dt \pi + (1 - \Lambda dt) w^{SS}].$$

Therefore the joint value of cooperation is

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

which is positive under Assumption 2. Comparing this with expression (2), 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 $\frac{\Lambda}{\Lambda + r}$ is strictly increasing in Λ , all-firm cooperation is welfare improving over any subset of firms' cooperation, including monopoly as a special case.

3.2.3 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 the two firms work on the arm until a discovery is made, the HJB equation gives us the following intuitive recursion:

$$w_n^{SS} = -\lambda_n c dt + e^{-rdt} [\lambda_n dt \pi + (1 - \lambda_n dt - \lambda_{-n} dt) w_n^{SS}], \quad (3)$$

where the third term is the discounted continuation payoff upon no discovery by either firm n or $-n$. The HJB equation immediately gives us

$$w_n^{SS} = \frac{\lambda_n}{\Lambda + r} (\pi - c). \quad (4)$$

Comparing with the single-firm case (2), 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 known return, competition is efficient.*

4 Equilibrium Analysis of the Model

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

4.1 Monopoly

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

$$w_n^R = -\lambda_n c dt + e^{-rdt} [\lambda_n dt \mu^0 \Pi + (1 - \lambda_n dt) w_n^R].$$

Note that there is no belief updating in the monopolistic problem. Hence $w_n^R = \frac{\lambda_n}{\lambda_n + r} (\mu^0 \Pi - c)$. 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 HJB equation,

$$v_n = -\lambda_n c dt + e^{-rdt} [\lambda_n dt (\mu^0 \Pi + w_n^S) + (1 - \lambda_n dt) v_n], \quad (5)$$

where $\mu^0 \Pi + w_n^S$ on the right-hand side 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. The HJB equation immediately gives us

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

This expression 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 $\lambda_n + r$.

4.2 Cooperation: Planner's Problem

We now consider the case in which firms behave cooperatively to maximize joint value. Several observations are in order.

1. Firms should share all the information to avoid wasteful research efforts.
2. Let w^{SS} and w^{RR} be the joint value of the two firms if they work *only* on the safe arm, and *only* the risky arm, respectively. Using an argument similar to that in the previous section,

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

By Assumptions 1–2, 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 arm until a discovery is made (which is made public immediately) and then both switch to the safe arm. Splitting the task is never optimal.

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

$$V = w^{RR} + \frac{\Lambda}{\Lambda + r} w^{SS}, \quad (6)$$

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

$$V_n = \frac{\lambda_n}{\Lambda} w^{RR} + \frac{\Lambda}{\Lambda + r} w_n^{SS}. \quad (7)$$

Proof. See 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 arm and then switching to the safe arm upon discovery, this also adds the continuation value of the safe research on top of w^{RR} . A discovery on the risky arm arrives at the rate Λ and the firm’s continuation payoff from the safe research upon arrival is simply w^{SS} .

Remark In fact, the proof of Proposition 2 shows that the strategy is optimal even without the restriction that both firms start with the risky arm.

4.3 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 now highlight how the ingredients in our model affect the learning dynamics.

First, we model two types of outcomes because such a model is more applicable to the prevalence of trial-and-error types of research competition.¹⁵ Uncertainty about the type of an opponent’s discovery is crucial for our learning dynamics generated by the dead-end discoveries.

Second, the independence of the arrival rates in the binary states implies that there will be no belief updating if research activities and dead-end findings are public. As a result, non-trivial 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 trade-off in the applications of our main interest.¹⁶

Third, arrival on the safe arm is also stochastic, which affects the learning dynamics indirectly. Upon observing an opponent’s discovery on the safe arm, a firm can make an inference about the opponent’s potential past observations on the risky arm, and the extent of this inference in equilibrium turns out to depend crucially on the timing of the safe arm 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 arm is ended.

We shall now demonstrate how learning takes place.

4.3.1 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 his opponent, firm $-n$, being on the risky arm at time t . Denote by b_n^t the probability that firm n assigns to his opponent being on the risky arm at time t conditional on the fact that the risky arm is bad.

Suppose both firms start on the risky arm, and switch only upon an observation. If firm n does not observe anything – neither from itself, nor from its opponent – from t

¹⁵The existing exponential bandit models have only a single outcome, for instance, Keller, Rady, and Cripps (2005), Bonatti and Hörner (2011), Strulovici (2010), Murto and Välimäki (2011), and Klein and Rady (2010).

¹⁶We discuss the relaxation of this assumption in the conclusion.

to $t + dt$, firm n will update μ_n^t using Bayes' rule as follows:

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

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 outcomes (see the discussion above).

The interpretation for the second equality above is as follows. The numerator measures the probability that the opponent does not make a (public) discovery and the risky arm 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 arm, 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 above Bayesian updating, we derive the law of motion for private beliefs¹⁷:

$$\dot{\mu}_n^t = -\mu_n^t (1 - \mu_n^t) b_n^t \lambda_{-n}. \quad (8)$$

The critical feature of the learning is that when the opponent discovers faster, i.e., 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 arm 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,

¹⁷To see this, subtract μ_i^t from both sides of the Bayes' formula, divide them by dt and then take the limits.

then

$$\beta_n^t = \frac{e^{-\lambda_{-n}t}}{e^{-\lambda_{-n}t} + (1 - \mu^0) \lambda_{-n}t e^{-\lambda_{-n}t}} = \frac{1}{1 + (1 - \mu^0) \lambda_{-n}t}. \quad (9)$$

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 ; $(1 - \mu^0) \lambda_{-n}t e^{-\lambda_{-n}t}$ is the probability that the opponent makes one dead-end discovery and that is the only discovery by time t .¹⁸ The denominator in (9) 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, $(1 - \mu^0) \lambda_{-n}t e^{-\lambda_{-n}t}$. The opponent will stay on the risky arm only when there is no arrival by $t \leq T$. This is reflected in the numerator of (9).

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,

$$b_n^t = \frac{e^{-\lambda_{-n}t}}{e^{-\lambda_{-n}t} + \lambda_{-n}t e^{-\lambda_{-n}t}} = \frac{1}{1 + \lambda_{-n}t}. \quad (10)$$

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 (9).

Substituting equation (10) into the filtering equation (8), we obtain

$$\dot{\mu}_n^t = -\mu_n^t (1 - \mu_n^t) \frac{\lambda_{-n}}{1 + \lambda_{-n}t}. \quad (11)$$

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, $\frac{\lambda_{-n}}{1 + \lambda_{-n}t}$, changes hyperbolically in time. The following lemma provides the explicit form for the belief.

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

$$\mu_n^t = \frac{\mu^0}{1 + (1 - \mu^0) \lambda_{-n}t}. \quad (12)$$

Proof. See Appendix. ■

Now, consider the case in which firm n has not discovered anything on its own

¹⁸Since the arrival rate is λ_{-i} , the probability of *one and only one* arrival by time t is

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

research, but observes the opponent's discovery on the safe research at $t \leq T$. Given the stopping strategy firm $-n$ adopts, firm n could infer that the opponent has already discovered a dead end on 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 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:

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

Finally, if firm n has not discovered anything on its own research at $t > T$, its belief μ_n^t is still μ_n^T . Note that there is a 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.

With the above discussion as a precursor, Figures 1A and 1B 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 arm, firm 1 stays on the risky arm until $T > 0$, and firm 2 sticks to the risky arm.¹⁹




FIGURE 1A: b_n^t




FIGURE 1B: μ_n^t

4.3.2 Equilibrium

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

¹⁹The parameters come from a numerical example provided in section 5.

Theorem 1 *Under Assumptions 1–2, there is a pure strategy perfect Bayesian equilibrium in which both firms start on the risky research and switch silently to the safe arm 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, i.e., $\frac{\mu^0 \Pi}{\pi}$ and $\frac{\lambda_2}{\lambda_1}$ are large enough, then the above describes the unique pure strategy equilibrium outcome.

Proof. See Appendix. ■

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 arm. 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 arm 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 arm. Moreover, the weak firm generates inefficiency from the time it switches until the strong firm discovers an outcome in the risky arm, due to early switching. In short, the weak firm endures two kinds of inefficiencies: *early-switching and dead-end inefficiencies*, while the larger firm endures only the *dead-end inefficiency*.

We also want to comment on the role of asymmetry. If firms are symmetric or payoffs in both arms are close, we can construct an equilibrium where firms coordinate on who switches research arms, and mixed strategy equilibria are also possible.

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

Proposition 3 *The equilibrium stopping time T is increasing in μ^0 and Π , and decreasing in λ_2 and π .*

Proof. See Appendix. ■

These comparative statics are intuitive. As μ^0 and Π become larger and π becomes smaller, the risky arm 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 non-monotonic as it affects both the weak firm’s payoffs in both arms simultaneously.

5 Numerical Example

In this section, we provide a numerical analysis of our model, taking pharmaceutical research competition as an example. Due to the simplicity of the model, our purpose is not to provide a detailed calibration of it. Our goal is rather to demonstrate the behaviour and welfare implications of the model, and highlight its general quantitative features for reasonable parameter values. Our model has 7 parameters: r , μ^0 , Π , π , c , λ_1 and λ_2 . Table 2 summarizes the parameter values. These parameters come from a basic calibration exercise in which we rely on reports by the Pharmaceutical Research and Manufacturers of America (*PhRMA*, 2002-2011). The details of the parameter choices are described in Appendix D.

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 2

5.1 Summary Statistics

Given the above parameters, we simulate the decentralized market and planner’s problem 500,000 times each. Table 3 summarizes the simulation results.

Both firms start on the risky arm with an initial belief $\mu_n^0 = 1/6$. As time elapses, firms receive outcomes according to the Poisson process described above. Note that firm 2 observes an outcome 2.5 times more frequently than the weak firm 1 (λ_2/λ_1). Since firm 2 receives an outcome faster, its average experimentation time on the risky arm 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 cut-off rule according to which it switches to the safe arm at $T = 36$ if it does not observe an outcome either by itself or from its competitor. The associated beliefs under this strategy were already depicted in Figures 1A and 1B.

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

TABLE 3

Figure 2A depicts the distribution for experimentation durations on the risky arm 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 cut-off strategy, switching to the safe arm 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 arm 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 hence reduces the average time spent on the risky arm to 10.9 months, which is 32% and 21% lower than the average experimentation times for firms 1 and 2, respectively.

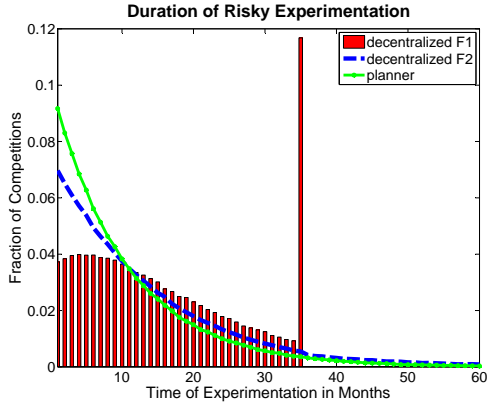


FIGURE 2A




FIGURE 2B

Next, we study the time that firms spend on risky research between two consecutive risky drug inventions. Figure 2B plots the results of the 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 arm 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 11 years -that is 26% less time- on the risky arm per drug.

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 arm before T , it switches to the safe arm 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 3A 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 cut-off strategy, which mitigates the welfare loss (however, as will be shown below, this strategy increases the second type of inefficiency). Since firm 2 learns the true nature of the arm faster, firm 1 spends more time on a dead-end risky arm before T . On the other hand, 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 for firm 2.

Figure 3B describes the second source of inefficiency: *early switching*. The planner prefers both firms to experiment until an outcome is found on the risky arm. 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 arm, as time

elapses. In equilibrium, firm 1 switches to the safe arm at time T even in situations where firm 2 has not received any information about the risky arm by then. This generates missing experimentations by firm 1 due to early switching, which are plotted in Figure 3B.

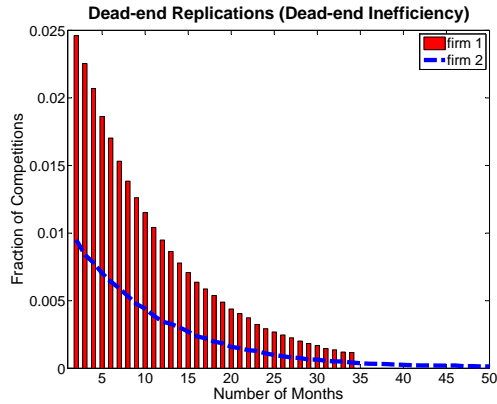


FIGURE 3A




FIGURE 3B

Finally, we illustrate the monetary cost of the problem in Figure 4, 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 inefficiencies. Firms spend on average \$382 million in the planner's economy, which is 23% less.




FIGURE 4

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

5.2 Types of Inefficiencies: Dead End and Early Switching

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 $W_{FB} = w^{RR} + \frac{\Lambda}{\Lambda+r}w^{SS}$. Since the intermediate regime differs from the first-best regime only after T , we have

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

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 arm and the safe arm, 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 arm 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 $\Pr(\text{one arrival before } t) = \lambda_n t e^{-\lambda_n t}$. Using this fact, we obtain²⁰

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

The following table summarizes the numerics. Note that firms do not want to share the dead-end discovery on the risky arm because of the competition on the safe arm,

²⁰This follows from: $W_I - W_D = (1 - \mu^0) \left\{ \int_0^T \lambda_2 t e^{-\lambda_2 t} e^{-(r+\lambda_1)t} \lambda_1 \pi dt + \int_0^T \lambda_1 t e^{-\lambda_1 t} e^{-(r+\lambda_2)t} \lambda_2 \pi dt \right. \\ \left. - \int_T^\infty \lambda_1 T e^{-\lambda_1 T} e^{-(r+\lambda_2)t} [e^{-\lambda_1(t-T)} \lambda_2 \pi + (1 - e^{-\lambda_1(t-T)}) \lambda_2 c] dt \right\}$.

which has a per unit of arrival rate net return $\pi - c$.

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	Early-switching 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%

TABLE 4

The finding is striking. We notice that even if the net return on the safe arm 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 the incentives to share private information. Without it, the firm does not lose anything from information sharing.

Note that the dead-end inefficiency is much larger than the early-switching inefficiency. We should not be optimistic about the early-switching inefficiency. Indeed, early-switching delays the discovery on the risky arm by almost 4 years for the same set of parameters as we demonstrated previously. If consumers' welfare is taken into account in our model, then early-switching will have a much larger implication.

6 Mechanism Design for Efficient Competition

In this section, we shall discuss a mechanism that incentivizes information sharing. 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., 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). 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.

Remark 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 has no incentive to pay anymore. This problem discourages information trading in a decentralized market. Therefore, a mediator is often necessary for the sale of information.

6.1 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:

- budget balance,
- a firm at *any point in time* should be allowed to walk away from the mechanism. That is, we face a design problem in which firms cannot commit to their future actions,
- 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
- 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 continuation decentralized 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 actions 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 behaviour 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 they 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.

6.2 Incentives

Henceforth we shall restrict our attention to a constant price path such that $p_n^t = p_n$.

6.2.1 No-delay Condition

Suppose 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,

$$\int_t^{t+h} e^{-(\Lambda+r)(\tau-t)} [-\lambda_n c + \lambda_n (\pi + p_n) + \lambda_{-n} (w_n^{ss} - p_{-n})] d\tau \leq p_n + w_n^{SS} \quad (13)$$

holds for any $h > 0$. Since $p_n \geq 0$, the RHS of (13) is strictly positive. Therefore, whenever the integrand in the LHS is negative, then (13) holds trivially. If the integrand

is strictly positive, the LHS is strictly increasing in h . Therefore, that (13) holds for any h is equivalent to

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

If instead $-\lambda_n c + \lambda_n (\pi + p_n) + \lambda_{-n} (w_n^{SS} - p_{-n}) > 0$, then since the LHS of (13) is increasing in h , (13) is equivalent to

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

This can be simplified into

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

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 arm: $\lambda_n (\pi - c - w_n^{SS})$.

Substituting w_n^{SS} into the above expression and simplifying, we have

$$\frac{\lambda_{-n} \lambda_n}{\Lambda + r} (\pi - c) \leq (\lambda_{-n} + r) p_n + \lambda_{-n} p_{-n}. \quad (14)$$

6.2.2 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 arm until firm $-n$ switches to the safe arm. Note that for firm 1, $v_{1,t}^S = v_{1,0}^S$ because firm 2 will never switch before a discovery. Therefore,

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

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. Hence

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

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,

$$p_1 \geq \frac{\lambda_1 \lambda_2}{(\Lambda + r)^2} (\pi - c) \quad \text{and} \quad p_2 \geq [1 - e^{-(\Lambda+r)T}] \frac{\lambda_1 \lambda_2}{(\Lambda + r)^2} (\pi - c). \quad (15)$$

6.2.3 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

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

The left-hand side is always V_n^D since 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

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

By Proposition 2, $\frac{\lambda_n}{\Lambda+r} (\mu^0 \Pi - c) + \frac{\Lambda}{\Lambda+r} w_n^{SS}$ on the right-hand side is firm n 's payoff V_n under full information sharing. Therefore, the condition can be rewritten as

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

This expression is very intuitive. The left-hand side is the expected net transfer firm n receives from participating in the mechanism: there will be transfer only when the risky arm 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 arm is $\frac{\lambda_n p_n - \lambda_{-n} p_{-n}}{\Lambda + r}$. The right-hand side 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.

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

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

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

where

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

It is feasible only when $\underline{K} \leq \overline{K}$. This condition is equivalent to

$$V_1^D + V_2^D \leq V_1 + V_2.$$

The right-hand side is the first-best joint payoff under full information. The left-hand side is the sum of values of the firms in the decentralized economy. Clearly, this condition is always satisfied.

6.3 Efficient Mechanism

Now, we summarize the two conditions on the prices:

1. No-delay condition:

$$\frac{\lambda_{-n} \lambda_n}{\Lambda + r} (\pi - c) \leq (\lambda_{-n} + r) p_n + \lambda_{-n} p_{-n}, \quad \text{for } n = 1, 2. \quad (16)$$

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

$$p_1 \geq \frac{\lambda_1 \lambda_2}{(\Lambda + r)^2} (\pi - c) \quad \text{and} \quad p_2 \geq [1 - e^{-(\Lambda + r)T}] \frac{\lambda_1 \lambda_2}{(\Lambda + r)^2} (\pi - c). \quad (17)$$

3. Participation constraint:

$$\underline{K} \leq \lambda_1 p_1 - \lambda_2 p_2 \leq \overline{K}. \quad (18)$$

Theorem 2 *Each price vector (p_1, p_2) that satisfies conditions (16) and (18) 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 (16)-(17) is non-empty. Indeed, we can set $p_1 = \frac{\lambda_2 p_2 + \bar{K}}{\lambda_1}$, which satisfies (18). 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 (16)-(17). 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.

6.4 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:

$$\min_{(p_1, p_2)} \{ \lambda_1 p_1 + \lambda_2 p_2 \} \quad \text{subject to} \quad \left\{ \begin{array}{l} \text{C1: } \frac{\lambda_1 \lambda_2}{\Lambda + r} (\pi - c) \leq (\lambda_1 + r) p_2 + \lambda_1 p_1 \\ \text{C2: } \frac{\lambda_1 \lambda_2}{\Lambda + r} (\pi - c) \leq (\lambda_2 + r) p_1 + \lambda_2 p_2 \\ \text{C3: } \frac{\lambda_1 \lambda_2}{(\Lambda + r)^2} (\pi - c) \leq p_1 \\ \text{C4: } [1 - e^{-(\Lambda + r)T}] \frac{\lambda_1 \lambda_2}{(\Lambda + r)^2} (\pi - c) \leq p_2 \\ \text{C5: } \underline{K} \leq \lambda_1 p_1 - \lambda_2 p_2 \leq \bar{K}. \end{array} \right.$$

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.

MINIMUM PRICE MECHANISM				
$\pi - c$	p_1^*	p_2^*	$\lambda_1 p_1^* + \lambda_2 p_2^*$	welfare recovery
\$ 1	\$ 0.5 (50¢)	\$ 0.20 (20¢)	\$ 0.02 (2¢)	\$19.3 million
\$ 1 million	\$ 0.5 million	\$ 0.2 million	\$ 0.02 million	\$19.6 million
\$ 10 million	\$ 4.7 million	\$ 1.8 million	\$ 0.24 million	\$22 million

TABLE 5

In the numerical computations, the two binding constraints of the mechanism are the no-delay condition for firm 1 (C1) and no-walk-away condition for firm 2 (C4).

The following graph plots the prices dictated by the minimum transfer mechanism as a function of the competition level on the safe research line.




FIGURE 5

Two features stand out in the above plot. 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 arm. For firm 1 this entails a larger reduction in value because it will then face a stronger competitor (firm 2).

7 Concluding Discussion and Extensions

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-arm bandit model with two asymmetric firms. We identified two types of inefficiencies that arise in this model and show 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 dead-end discovery and asymmetric information. In what follows, we shall discuss possible extensions of our model and future research.

7.1 State-dependent Arrival Rate

In this paper, we have assumed that the arrival rate λ_n is independent of the true state. More generally, one might allow the arrival rate to be a function of the state as well, λ_n^s , where $s \in \{G, B\}$ where G stands for the good risky arm and B stands for the dead-end risky arm. A source of *exogenous* learning shows up in this environment. For instance,

if $\lambda_n^G \neq \lambda_n^B$, then firm n will learn from the fact that there is no discovery from its own research. In particular, if $\lambda_n^G > \lambda_n^B$, then for firm n , no news from its own research is bad news. In this case, learning from n 's own research and learning from the opponent's research (no discovery) reinforce each other. If, instead, $\lambda_n^G < \lambda_n^B$, then no news is good news. Therefore, learning from n 's own research and learning from the opponent tend to push the learning in different directions. Our model isolates the endogenous learning through competition from the exogenous learning. It remains to analyze which force will be stronger and how they interact over time. We believe this complication will not change the qualitative predictions of our model.

7.2 Strategic Patenting

In our model, a firm receives a lump-sum payoff from its good discovery immediately. We could enrich the model to study strategic patenting decisions and ask whether a firm has an incentive to delay its patenting decision to its own benefit. In this section, we shall argue that the equilibrium we characterize is robust to an endogenous patenting decision. Therefore, to study strategic patenting decisions, we need to enrich the model (for example, by allowing multiple arrivals). This is an interesting question to ask but is orthogonal to the current focus.

Assume firm $-n$'s strategy is to patent its discovery immediately. Consider firm n . If firm n has a non-patented successful discovery at a point when the other firm has already switched, then there is no benefit from delayed patenting, and a cost due to discounting. Now consider the case where firm n has a non-patented discovery at t when the competing firm is still working on the risky research (note that such a discovery may be made exactly at t or it is discovered before but delayed until t). If n patents this discovery at t , then we can derive its payoff as $V_t = \Pi + w_n^{SS}$. Suppose the firm decides to delay it until $t + s$, for some $s > 0$. Since we know that the firm will not delay patenting when the other firm has switched, we can assume without loss of generality that at $t + s$ firm $-n$ is still on the risky arm. Firm n 's expected payoff at t is therefore,

$$V_{t,s} = \int_t^{t+s} e^{-(r+\Lambda)(\tau-t)} [\lambda_n (\pi + \Pi) + \lambda_{-n} w_n^{SS} - \lambda_n c] d\tau + e^{-(r+\Lambda)s} (\Pi + w_n^{SS}).$$

Now

$$\begin{aligned}
\frac{\partial V_{t,s}}{\partial s} &= e^{-(r+\Lambda)s} [\lambda_n (\pi + \Pi) + \lambda_{-n} w_n^{SS} - \lambda_n c] - (r + \Lambda) e^{-(r+\Lambda)s} (\Pi + w_n^{SS}) \\
&= e^{-(r+\Lambda)s} \left\{ -r\Pi - \lambda_{-n} \left[\Pi - \frac{\lambda_n (\pi - c)}{r + \Lambda} \right] \right\} \\
&< e^{-(r+\Lambda)s} [-r\Pi - \lambda_{-n} (\Pi - \pi)].
\end{aligned}$$

Note that under Assumption 1, $\Pi > \pi$ and hence $\frac{\partial V_{t,s}}{\partial s} < 0$. Therefore, firm n , if it has a non-patented innovation at time t , will not delay patenting by any $s > 0$.

Remark 1 *Note that we have just shown that it is optimal for firm n to patent immediately when firm $-n$'s strategy is registering immediately whenever $\Pi > \pi$. The intuition is that if firm n delays for dt , the cost of delay is of the order $\lambda_{-n}\Pi$, yet the benefit is $\lambda_{-n}\pi$ because firm n keeps firm $-n$ away for dt .*

7.3 Macroeconomic Applications

The increase in potentially wasteful R&D dollars has been a common concern both in academic and policy spheres. Macro data on innovation and R&D spending in the US exhibits a worrisome time-series pattern. The ratio of registered innovation counts to total innovation efforts in the US has been steadily decreasing over time. The following two figures document this stylized fact.




FIGURE 6A




FIGURE 6B

In figure 6A, we plot the ratio of the total number of USPTO patents granted to US residents over aggregate R&D investment in the US. In the early 1950s, the patent-R&D ratio was around 1.4 and it has decreased by almost 70% to 0.4 in the early 2000s. There could be various explanations for this decline, and Kortum (1993) argues that one of

them is the increasing duplicative R&D efforts by competing firms. He suggests that the increase in market size leads to a larger ex-post value of innovation, which, combined with competition, leads to a larger R&D spending per patent. A similar and even more drastic picture emerges in the pharmaceutical industry. Figure 6B plots the number of drug approvals per R&D investment for this industry. The ratio declines from 1.4 in the early 1960s to 0.1 in the early 2000s, which is a decline of more than a 90%. This observation again hints at a severe problem of R&D duplication for drug inventions.

We provide two comparative statics as a preliminary attempt to use our model to touch on this issue. The first one is the increase in the market value of drugs. Although this increase in value could be caused by many different factors (increase in market size, for instance), the end effect is an increase in the ex-post returns to innovation. In our model, an increase in the market value of drugs leads to more experimentation on the risky arm, which causes an increase in the cut-off value T of the weak firm 1. This in turn also increases dead-end replications and reduces the number of drugs per R&D investment. Figure 7A plots the average number of drugs per R&D investment as a function of the market value.




FIGURE 7A




FIGURE 7B

Another potential explanation emerging from our model is the increase in uncertainty or the decrease in the probability of a good outcome on the risky arm.²¹ To understand this better, consider the case in which $\mu^0 = 1$ and the decentralized equilibrium would be efficient. As uncertainty increases (μ^0 declines), the decentralized economy increases the amount of dead-end replications. This increase in wasteful spending reduces the number of drugs per R&D investment as illustrated in Figure 7B.

A more detailed analysis of the macroeconomic implications of the inefficiencies identified in this paper requires incorporating the microeconomic structure into a formal

²¹In reality this could be caused due to the fishing-out effect.

general equilibrium growth model. Akcigit and Liu (2011) take a step in this direction. We believe that additional interesting macroeconomic questions are still awaiting future exploration.

8 References

- ACEMOGLU, D., AND U. AKCIGIT (2011) “Intellectual Property Rights Policy, Competition and Innovation”, *Journal of the European Economic Association*, forthcoming.
- ACEMOGLU, D., K. BIMPIKIS, AND A. OZDAGLAR (2011) “Experimentation, Patents and Innovation”, *American Economic Journal: Microeconomics*, 3, 37-77.
- AGHION, P., C. HARRIS, P., AND J. VICKERS (2001) “Competition, Imitation and Growth with Step-by-Step Innovation”, *Review of Economic Studies*, 68, 467-492.
- AGHION, P., AND P. HOWITT (1992) “A Model of Growth Through Creative Destruction”, *Econometrica*, 110, 323-351.
- AKCIGIT U., AND Q. LIU (2011) “Endogenous Growth with Asymmetric Information: A Mechanism Design Approach”, Mimeo, University of Pennsylvania.
- ARROW, K. (1962) “Economic Welfare and the Allocation of Resources for Invention”, in R.R. Nelson, ed. Universities-National Bureau of Economic Research Conference Series; *The Rate and Direction of Economic Activities: Economic and Social Factors*, Princeton University Press, New York.
- BERGEMANN, D., AND J. VÄLIMÄKI (2008) “Bandit Problems”, in Durlauf, S. and Blume, L. (eds.) *The New Palgrave Dictionary of Economics*, 2nd edn (Basingstoke: Palgrave Macmillan Ltd).
- BHATTACHARYA, S., K. CHATTERJEE, AND L. SAMUELSON (1986) “Sequential Research and the Adoption of Innovations”, *Oxford Economic Papers*, 38, 219–243.
- BOLTON, P., AND C. HARRIS (1999) “Strategic Experimentation”, *Econometrica* 67, 349–74.
- BONATTI, A., AND J. HÖRNER (2011) “Collaborating”, *American Economic Review*, 101, 632-663.

- CHATTERJEE, K., AND R. EVANS (2004) “Rivals Search for Buried Treasure: Competition and Duplication in R&D”, *RAND Journal of Economics*, 15, 160-183.
- FERSHTMAN, C., AND A. RUBINSTEIN (1997) “A Simple Model of Equilibrium in Search Procedures”, *Journal of Economic Theory*, 72, 432-441.
- GRABOWSKI, H., J. VERNON, AND J. A. DIMASI (2002) “Returns on Research and Development for 1990s New Drug Introductions,” *Pharmacoeconomics*, 20, 11-29.
- HARRIS, C., AND J. VICKERS (1985) “Perfect Equilibrium in a Model of Race”, *Review of Economic Studies*, 52, 193-209.
- _____ (1987) “Racing with Uncertainty”, *Review of Economic Studies*, 54, 1-21.
- HOPENHAYN, H., G. LLOBET, AND M. MITCHELL (2006) “Rewarding Sequential Innovators: Prizes, Patents and Buyouts”, *Journal of Political Economy*, 114, 1041-1068.
- HOPENHAYN, H., AND M. MITCHELL (2011) “Optimal Patent Policy with Recurrent Innovators”, UCLA and University of Toronto Working Paper.
- HOPENHAYN, H., AND F. SQUINTANI (2011) “Preemption Games with Private Information”, *Review of Economic Studies*, 78, 667-692.
- KELLER, G., S. RADY, AND M. W. CRIPPS (2005) “Strategic Experimentation with Exponential Bandits”, *Econometrica* 73, 39–68.
- KLEIN, N., AND S. RADY (2011) “Negatively Correlated Bandits”, *Review of Economic Studies*, 78, 693-732.
- KORTUM, S. (1993) “Equilibrium R&D and the Patent-R&D Ratio: US Evidence,” *American Economic Review Papers and Proceedings*, 83, 450-457.
- KREMER, M. (1998) “Patent Buy-outs: A Mechanism for Encouraging Innovation”, *Quarterly Journal of Economics*, 113, 1137–1167.
- MOSCARINI, G., AND L. SMITH (2001) “The Optimal Level of Experimentation”, *Econometrica*, 69, 1629–1644.

- MOSCARINI, G., AND F. SQUINTANI (2010) “Competitive Experimentation with Private Information: The Survivor’s Curse”, *Journal of Economic Theory*, 145, 639–660.
- MURTO, P., AND J. VÄLIMÄKI (2011) “Learning and Information Aggregation in an Exit Game”, *Review of Economic Studies*, forthcoming.
- PHRMA (2011) “Pharmaceutical Industry: Profile 2011”, The Pharmaceutical Research and Manufacturers of America, Washington, D.C.
- REINGANUM, J. (1982) “A Dynamic Game of R and D: Patent Protection and Competitive Behavior”, *Econometrica*, 671–688.
- ROSENBERG, D., E. SOLAN, AND N. VIEILLE (2007) “Social Learning in One-Armed Bandit Problems”, *Econometrica*, 75, 1591–1611.
- SCOTCHMER, S. (1991) “Standing on the Shoulders of Giants: Cumulative Research and the Patent Law”, *Journal of Economic Perspectives*, 5, 29–41.
- STRULOVICI, B. (2010) “Learning While Voting: Determinants of Collective Experimentation”, *Econometrica*, 78, 933–971.
- THOMAS C. (2011) “Experimentation with Congestion”, Mimeo, University of Texas at Austin.
- TUFTS CENTER FOR THE STUDY OF DRUG DEVELOPMENT (2009) “Outlook 2009”, Tufts University.

Appendix

A Proof of Proposition 2

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

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

which implies

$$V = \frac{\Lambda}{\Lambda + r} (\mu^0 \Pi - c + w^{SS}). \quad (19)$$

This joint value can be also rewritten as

$$\begin{aligned} V &= \frac{\Lambda (\mu^0 \Pi - c)}{\Lambda + r} + \frac{\Lambda}{\Lambda + r} \frac{\Lambda (\pi - c)}{\Lambda + r} \\ &= w^{RR} + \frac{\Lambda}{\Lambda + r} w^{SS}. \end{aligned} \quad (20)$$

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

Proof of Proposition 2. 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 HJB equation,

$$\hat{V} = \max \left\{ \begin{array}{l} \Lambda dt (\mu^0 \Pi + w^{SS}) e^{-rdt} - \Lambda c dt + (1 - \Lambda dt) \hat{V} e^{-rdt}, \\ \Lambda dt (\pi + w^{RR}) e^{-rdt} - \Lambda c dt + (1 - \Lambda dt) \hat{V} e^{-rdt}, \\ \lambda_1 dt (\mu^0 \Pi + w^{SS}) e^{-rdt} + \lambda_2 dt (\pi + w^{RR}) e^{-rdt} - \Lambda c dt + (1 - \Lambda dt) \hat{V} e^{-rdt}, \\ \lambda_2 dt (\mu^0 \Pi + w^{SS}) e^{-rdt} + \lambda_1 dt (\pi + w^{RR}) e^{-rdt} - \Lambda c dt + (1 - \Lambda dt) \hat{V} e^{-rdt} \end{array} \right\} \quad (21)$$

where the four terms on the right side are the payoffs from strategies in which both firms start with the risky arm, both firms start with the safe arm, firm 1 starts with the risky arm and firm 2 starts with the safe arm, and firm 2 starts with the risky arm and firm 1 starts with the safe arm, respectively.

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

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

Note that the inequality follows from Assumption 1. Therefore, the first term on the right side of (21) 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 arm. 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 (20),

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

This completes the proof. ■

B 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 (11) we get

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

which reduces to

$$B + B\lambda_{-n}t = (1 - A)\lambda_{-n} + \lambda_{-n}Bt.$$

Equating the constant terms we get $B = (1 - A)\lambda_{-n}$. Moreover, we impose the boundary condition $\Psi(0) = \mu^0$. Then we get $A = \mu^0$ and $B = (1 - \mu^0)\lambda_{-n}$. This verifies our conjecture. ■

C Proofs of Theorem 1 and Proposition 3

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 arm is μ_n^t and his belief on its opponent, firm $-n$, is still on the risky arm is β_n^t . Recall from Equation (4) that w_n^{SS} is firm n 's expected payoff from competing with firm $-n$ on the safe arm, $w_n^{SS} = \frac{\lambda_n}{\Lambda+r} (\pi - c)$.

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

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

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

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

The first line is firm 1's discounted expected return when it makes a discovery on the risky arm 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 discovers. It again depends on the position of firm 2. If firm 2 is on the risky arm, which happens with probability β_1^t , firm 1 will compete with firm 2. If firm 2 is on the safe arm, a discovery on the safe arm indicates the risky arm 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 arm and staying there forever is given by

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

The interpretation is similar to the previous case.

Therefore, by taking the limit, the indifference condition becomes

$$\mu_1^t (\Pi + w_1^{SS}) + (1 - \mu_1^t) [b_1^t v_1^S + (1 - b_1^t) w_1^{SS}] = \pi. \quad (22)$$

This condition carries the following intuition. At time t , spending an additional amount of time dt on either arm 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 arm. The LHS is the expected return on the risky arm. With probability μ_1^t , the arm is good, in which case firm 1 receives the patent value Π and competes with firm 2 in the safe arm and obtains w_1^{SS} . With the remaining probability $(1 - \mu_1^t)$ the arm is bad, in which case, firm 1 switches secretly to the safe arm and obtains a payoff depending on whether firm 2 is already on the safe arm or not.

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

$$\mu_1^T \Pi + (1 - \mu_1^T) b_1^T (v_1^S - w_1^{SS}) + w_1^{SS} = \pi \quad (23)$$

From equations (10) and (12), we know that for $n = 1, 2$,

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

Hence

$$\begin{aligned} T &= \underbrace{\frac{1}{(1 - \mu^0) \lambda_2}}_{\text{learning channel}} \underbrace{[\mu^0 (\Pi + w_1^{SS}) + (1 - \mu^0) v_1^S - \pi]}_{\text{Risky research premium}} \underbrace{\left(\frac{1}{\pi - w_1^{SS}} \right)}_{\text{Competition Channel}} \\ &= \frac{1}{(1 - \mu^0) \lambda_2} \left[\frac{\mu^0 \Pi + (1 - \mu^0) (v_1^S - w_1^{SS})}{(\pi - w_1^{SS})} - 1 \right] \\ &= \frac{1}{(1 - \mu^0) \lambda_2} \left[\frac{\mu^0 \Pi}{\pi} \frac{(r + \Lambda)}{(r + \Lambda - \lambda_1 \frac{\pi - c}{\pi})} - 1 \right] + \frac{\lambda_1 \frac{\pi - c}{\pi}}{(r + \Lambda) (r + \Lambda - \lambda_1 \frac{\pi - c}{\pi})} \end{aligned}$$

Remark 2 (Proposition 3) *From the explicit expression for T above, it is easy to check 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, given that both start on the risky research line. In Step 4, after we 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 firm 2 does not stop the risky research before a discovery. Recall that T is the unique solution of

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

That is, T uniquely solves

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

We know the LHS is monotone decreasing in t . Hence if $t < T$, firm 1 strictly prefers to stay in the risky arm, 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 found.

Now assume 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 arm in the candidate equilibrium is given by the recursion:

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

Note that since $\mu_1^T \Pi - c \geq 0$ (otherwise, firm 1 would have already switched to the safe arm before T), $\mu_2^T \Pi - c > 0$ by (12). Hence

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

In order for firm 2 to stay on the risky research, we need $V_2 \geq w_2^{SS}$. Plugging in param-

eters, the sufficient condition can be simplified progressively as

$$\begin{aligned}
-\lambda_2 c + \lambda_2 (\mu_2^T \Pi + w_2^{SS}) + \lambda_1 \frac{\lambda_2 (\mu_2^T \Pi - c)}{r + \lambda_2} &\geq (r + \Lambda) w_2^{SS} \\
\mu_2^T \Pi \left(1 + \frac{\lambda_1}{r + \lambda_2}\right) + w_2^{SS} - \frac{\lambda_1 c}{r + \lambda_2} &\geq \pi \\
\mu_2^T \Pi - \pi + w_2^{SS} + (\mu_2^T \Pi - c) \frac{\lambda_1}{r + \lambda_2} &\geq 0
\end{aligned} \tag{24}$$

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 arm is higher and the return on the safe arm 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 firm 1 does not face any competition in the risky arm but faces competition with certainty on the safe arm (continuing with the same belief updating). In that case the indifference condition in (23) 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 (24) is

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

Using the expression for μ_1^{T*} , the sufficient condition becomes

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

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 prescribed above. The reason is that the alternative strategy of firm 1 increases the competition of the risky arm and reduces the competition on the safe arm. 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 (23) holds:

$$\pi = \mu_2^{T_2} \Pi + (1 - \mu_2^{T_2}) b_2^{T_2} (v_2^S - w_2^{SS}) + w_2^{SS}. \quad (25)$$

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:

$$\begin{aligned} \mu_2^{T_2} &\geq \mu_2^T, \\ \mu_2^{T_1} &> \mu_1^T, \\ (1 - \mu_2^T) b_2^T &> (1 - \mu_1^T) b_1^T, \\ v_2^S - w_2^{SS} &> v_1^S - w_1^{SS}, \\ w_2^{SS} &> w_1^{SS}. \end{aligned}$$

Utilizing these inequalities, we derive from (25) that

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

A contradiction.

Step 3: (Uniqueness) There are no other equilibrium stopping strategies when $\frac{\lambda_2}{\lambda_1}$ and $\frac{\mu^0 \Pi}{\pi}$ 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: $+\infty \geq T_1 > T_2$.

We define $v_2^S(T_2, T_1)$ as the value of firm 2 at T_2 when it switches to the safe arm but anticipating that firm 1 might switch to the safe arm 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 for another dt ,” is

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

The payoff from “spend the next dt on the safe arm and stay there forever,” is given by

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

Therefore, by taking the limit, the indifference condition becomes

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

or, equivalently,

$$\mu_2^{T_2} \Pi + (1 - \mu_2^{T_2}) b_2^{T_2} [v_2^S (T_2, T_1) - w_2^{SS}] + w_2^{SS} = \pi. \quad (26)$$

Notice that $v_2^S (T_2, T_2) = w_2^{SS} \leq v_2^S (T_2, T_1)$ for any $T_1 > T_2$. Then (26) gives us

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

which is

$$T_2 \geq \frac{\mu^0 \Pi - (\pi - w_2^{SS})}{(\pi - w_2^{SS}) (1 - \mu^0) \lambda_1}. \quad (27)$$

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

i.e.,

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

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

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

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,

$$T_2 \leq \frac{1}{(1 - \mu^0) \lambda_2} \left[\frac{\mu^0 \Pi}{\frac{\lambda_2 + r}{\Lambda + r} \frac{\lambda_1 + r}{\Lambda + r} (\pi - c) + c} - 1 \right] \quad (28)$$

Comparing (27) and (28), a contradiction will be derived if

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

which is equivalent to

$$\mu^0 \Pi \left[\frac{\lambda_2}{\frac{\lambda_1 + r}{\Lambda + r} (\pi - c) + c} - \frac{\lambda_1}{\frac{\lambda_2 + r}{\Lambda + r} \frac{\lambda_1 + r}{\Lambda + r} (\pi - c) + c} \right] > \lambda_2 - \lambda_1. \quad (29)$$

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

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

Hence a sufficient condition for (29) 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: $+\infty > T_2 \geq T_1$. In this case, firm 2 does not update its belief after T_1 if it does not observe anything on the risky arm. 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 arm) and staying on the risky arm (monopolizing the risky arm with the option value of the safe arm) at any $t \geq T_1$. Following the argument in the previous case, the indifference condition of firm 1 is

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

Recall that our equilibrium indifference condition is given by

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

Since $b_n^T (1 - \mu_n^T) = \frac{1 - \mu^0}{1 + (1 - \mu^0) \lambda_{-n} T}$, the LHS of the previous equation is strictly decreasing in T . Now suppose $T \leq T_1$. Then it follows from $v_1^S > v_1^S(T_1, T_2)$ that

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

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

In our equilibrium, firm 2 prefers to stay on the risky arm 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

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

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 arm until a discovery or

switching at T_1). Hence

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

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

D Details of the Numerical Example

Our model has 7 structural parameters: r , μ^0 , Π , π , c , λ_1 and λ_2 . We take the annual interest rate to be $r = 5\%$. According to PhRMA (2011) only one out of six drug candidates survives the clinical stage; thus, we set $\mu^0 = 1/6$. The remaining five parameters are calibrated to the relevant moments from the data.

Our strategy is to calibrate the model to the clinical trial stage of the pharmaceutical research during the late 1990s. The analysis requires the empirical characterization of two asymmetric firms. For this purpose, we use the population of pharmaceutical companies in Compustat in 2000. Since the strength of the firms is determined by the R&D spendings in our model, we rank the firms in the Compustat sample according to their R&D investments in 2000. We form the strong firm by averaging the numbers of the top 3% companies in Compustat. Similarly, the weak firm is formed by averaging the second top 3% percent companies. The following table summarizes the empirical target moments and their data sources²²:

SOME KEY FACTS ON PHARMACEUTICAL R&D AND CALIBRATION TARGETS		
Moment Description	Data	Model
[‡] Average time to develop a drug	10-15 years	14.8 years
[‡] Fraction of candidate drugs that survive the clinical trial	1/6	1/6
*Net present value of a drug	\$1.4 billion	\$1.4 billion
*Average cost to develop a drug	\$480 million	\$496 million
[§] Ratio of R&D spendings	2.5	2.5
[§] Ratio of profits	2.5	2.9

^{22‡}Obtained from PhRMA (2011). *Obtained from Grabowski, Vernon and DiMasi (2002). [§]Obtained from Compustat (dnum=2834) for 2005. Ratios are defined as the strong firm's moment divided by the weak firm's moment. Profits are computed as: Revenue-R&D-Cost of goods sold. Rate of return to R&D is the ratio of profit to R&D.

Note that our calibrated model delivers a cutoff time $T = 36$, which means that the weak firm experiments in the risky research line for 36 months as long as it neither receives an outcome from its own research effort nor observes a patent from the competitor firm. As discussed in the main text, this is one of the key sources of inefficiency in this competition.

Supplementary Material for:
“The Role of Information in Competitive Experimentation”

Ufuk Akcigit

Qingmin Liu

November 8, 2011

In the main text, we have assumed that the game starts on the risky arm. This section considers a model in which both firms have no research activity before $t = 0$, and simultaneously, right at $t = 0$, each of them has to decide which arm to take to start the game. In particular, a firm can start on the safe arm and then switch to the risky arm, or it could choose not to research at all.

Theorem 3 *The equilibrium described in Theorem 1 in the main text is the unique pure strategy equilibrium when firms can choose the initial starting arm freely, provided that there is enough asymmetry across research lines and players, i.e., $\frac{\mu^0\Pi}{\pi}$ and $\frac{\lambda_2}{\lambda_1}$ are large enough.*

Proof Given the proof for Theorem 1, we need to show two additional claims. In Step 1 below, we shall show that there is no equilibrium in which either player starts with the safe research line. In Step 2 below, we verify that the initial choices of the risky arm are best responses to each other in the candidate equilibrium. In particular, we need to verify that the following deviation is unprofitable for a firm: start on the safe arm with a hope that it can make a discovery before T , which will fool the opponent into thinking that the risky arm had already been discovered to be a dead end; hence the opponent is misled into quitting the entire game, leaving the risky arm to the deviating firm. This deviation is not possible in our benchmark model as a firm cannot return to the risky arm there.

Step 1: We shall also show that there is no equilibrium in which either player starts with the safe research line if $\mu^0\Pi > \pi \left(1 + \frac{\Lambda}{r+\lambda-n} + \frac{\lambda-n}{\lambda_n+r}\right)$. There are several cases to consider.

Case 1.1: Both firms start on the safe arm, with stopping times $T_1, T_2 \in (0, +\infty]$, respectively. We claim that $T_1 = T_2 = T^* \in (0, +\infty]$. Suppose for the purpose of contradiction that $T_n > T_{-n}$, then upon no observation of discovery from T_{-n} on, firm n 's belief will become more pessimistic over time. Consequently, if firm n does not want to switch at T_{-n} upon no discovery, it will not switch at any future time upon no discovery. That is, $T_n = \infty$. Now, we have a situation in which firm n works on the safe arm until a discovery and firm $-n$ starts with the safe arm but switches at T_{-n} . Since firm $-n$'s belief on the risky arm will never get updated, the firm should instead start with the risky arm at $t = 0$. A contradiction. Hence the only possibility left is $T_n = T_{-n} > T^* \in (0, +\infty]$.

Then, firm n 's expected payoff will be

$$V_n = \int_0^{T^*} e^{-(\Lambda+r)t} [\lambda_n (\pi + v_n^{RR}) + \lambda_{-n} v_n^{RR} - \lambda_n c] dt + e^{-(\Lambda+r)T^*} v_n^{RR}, \quad (30)$$

where v_n^{RR} is firm n 's expected payoff of competing with firm $-n$ on the risky research with 0 outside options (because the outcome on the safe arm has been discovered).

Now fix firm $-n$'s strategy and consider a deviation of firm n of starting with the risky arm until a discovery by itself. Firm n 's payoff will be *at least*

$$V_n^d = \int_0^{T^*} e^{-(\Lambda+r)t} [\lambda_n (\mu_0 \Pi + w_n^{SS}) + \lambda_{-n} v_n^{RR} - \lambda_n c] dt + e^{-(\Lambda+r)T^*} v_n^{RR}. \quad (31)$$

The reason for V_n^d being a lower bound is that conditional on no discovery up to time T^* , the continuation payoff for firm n is *at least* v_n^{RR} because firm n still has the option of going to the safe arm.

Notice that

$$\begin{aligned} V_n^d - V_n &= \int_0^{T^*} e^{-(\Lambda+r)t} [\lambda_n (\mu_0 \Pi + w_n^{SS}) + \lambda_{-n} v_n^{RR} - \lambda_n (\pi + v_n^{RR}) - \lambda_{-n} v_n^{RR}] dt \\ &= \int_0^{T^*} e^{-(\Lambda+r)t} [\lambda_n (\mu_0 \Pi - \pi + w_n^{SS} - v_n^{RR})] dt. \end{aligned}$$

Since firms' total payoff with competition on the risky arm without the option of the safe arm is less than the cooperative counterpart, we have $v_n^{RR} + v_{-n}^{RR} < \frac{\lambda_n (\mu_0 \Pi - c)}{\Lambda + c} + \frac{\lambda_{-n} (\mu_0 \Pi - c)}{\Lambda + c}$.

Hence for at least one $n = 1, 2$, $v_n^{RR} < \frac{\lambda_n(\mu_0\Pi - c)}{\Lambda + c}$. For this n , we have

$$\begin{aligned} V_n^d - V_n &> \int_0^{T^*} e^{-(\Lambda+r)t} \left[\lambda_n \left(\mu_0\Pi - \pi + w_n^{SS} - \frac{\lambda_n(\mu_0\Pi - c)}{\Lambda + c} \right) \right] dt \\ &= (\mu_0\Pi - \pi) \lambda_n \frac{\lambda_{-n} + r}{\Lambda + r} \int_0^{T^*} e^{-(\Lambda+r)t} dt \\ &> 0. \end{aligned}$$

Hence, for this firm n , deviation is profitable.

Case 1.2: Firm n starts on the safe arm with stopping time $T_n \in (0, +\infty]$. Firm $-n$ starts on the risky arm, with stopping time $T_{-n,0} \in (0, +\infty]$, and $T_{-n,1} \geq 0$ (the second stopping time is for the stage in which firm n discovers on the safe arm).

Consider the subgame right after firm n takes the safe arm. We modify firm n 's problem as follows:

- (a) Fix firm $-n$'s strategy as staying on the risky arm forever until a discovery is observed on the risky arm. Let \tilde{T}_n be firm n 's one optimal stopping time in this auxiliary problem. We claim that in this auxiliary problem we can take $\tilde{T}_n > 0$. Indeed, $\tilde{T}_n \geq T_n$. The reason is that this modification makes staying on the safe arm for any $t > 0$ more attractive than in the original problem (the potential benefit from the risky arm is reduced while the benefit from the safe arm is increased because firm n will face less competition there).
- (b) On top of (a), ask firm $-n$ to reveal its discovery (including the dead-end finding) until firm n leaves the safe arm.²³ Hence, at any t , by which no discovery is made, there is no belief updating. Therefore, if firm n starts with the safe arm in the auxiliary problem (a), then it will always stay on the safe arm before a discovery.

Let V_n^{RR} be firm n 's payoff upon switching to the risky arm in the auxiliary problem (b). Firm n 's expected payoff at time 0 in this auxiliary problem can be written as

$$\begin{aligned} &\int_0^\infty e^{-(\Lambda+r)t} \left[\lambda_n (\pi + V_n^{RR} - c) + \lambda_{-n} w_n^{SS} \right] dt \\ &= \frac{\lambda_n}{\Lambda + r} V_n^{RR} + \frac{\lambda_n}{\Lambda + r} (\pi - c) \left(1 + \frac{\lambda_{-n}}{\Lambda + r} \right). \end{aligned}$$

²³Note that we construct this artificial problem for firm n where firm $-n$'s strategy is superimposed exogenously. This should not be confused with the observability assumption in the original problem.

Because firm $-n$'s strategy is exogenously fixed as in (a), V_n^{RR} is independent of π .

Consider an alternative strategy for firm n in the auxiliary problem (b): abandon the safe arm immediately. The expected payoff from this alternative strategy is V_n^{RR} . A contradiction arises if

$$\frac{\lambda_n}{\Lambda + r} V_n^{RR} + \frac{\lambda_n}{\Lambda + r} (\pi - c) \left(1 + \frac{\lambda_{-n}}{\Lambda + r} \right) < V_n^{RR}$$

which is equivalent to

$$\pi - c < V_n^{RR} \left(\frac{\lambda_{-n} + r}{\lambda_n} \frac{\Lambda + r}{\Lambda + r + \lambda_{-n}} \right).$$

Note that $V_n^{RR} \geq \mu^0 \frac{\lambda_n(\Pi - c)}{\Lambda + r} - (1 - \mu^0) \frac{\lambda_n c}{\lambda_n + r}$. Hence, a sufficient condition for the above expression is

$$\begin{aligned} \pi - c &< \left[\mu^0 \frac{\lambda_n(\Pi - c)}{\Lambda + r} - (1 - \mu^0) \frac{\lambda_n c}{\lambda_n + r} \right] \left(\frac{\lambda_{-n} + r}{\lambda_n} \frac{\Lambda + r}{\Lambda + r + \lambda_{-n}} \right) \\ &= \left[\mu^0 (\Pi - c) - (1 - \mu^0) c \frac{\Lambda + r}{\lambda_n + r} \right] \frac{\lambda_{-n} + r}{\Lambda + r + \lambda_{-n}} \\ &= \left[\mu^0 \Pi - c - (1 - \mu^0) c \frac{\lambda_{-n}}{\lambda_n + r} \right] \frac{\lambda_{-n} + r}{\Lambda + r + \lambda_{-n}}. \end{aligned}$$

This is

$$\begin{aligned} \mu^0 \Pi - c &> (\pi - c) \frac{\Lambda + r + \lambda_{-n}}{r + \lambda_{-n}} + (1 - \mu^0) c \frac{\lambda_{-n}}{\lambda_n + r} \\ \mu^0 \Pi - c &> \pi \frac{\Lambda + r + \lambda_{-n}}{r + \lambda_{-n}} - c \frac{\Lambda + r + \lambda_{-n}}{r + \lambda_{-n}} + (1 - \mu^0) c \frac{\lambda_{-n}}{\lambda_n + r} \\ \mu^0 \Pi &> \pi \frac{\Lambda + r + \lambda_{-n}}{r + \lambda_{-n}} - c \frac{\Lambda}{r + \lambda_{-n}} + (1 - \mu^0) c \frac{\lambda_{-n}}{\lambda_n + r} \end{aligned}$$

A sufficient condition is given by

$$\frac{\mu^0 \Pi}{\pi} > 1 + \frac{\Lambda}{\lambda_{-n}} + \frac{\lambda_{-n}}{\lambda_n}.$$

Therefore, under the above condition, working on the safe arm is not optimal.

Step 2: Best responses of the initial choices in the candidate equilibrium.

We shall utilize the idea of the auxiliary problems in Step 1. Suppose firm n has a

profitable deviation that consists of starting on the safe arm with stopping time $\tilde{T}_n \in (0, +\infty]$. Now in the auxiliary problem (a) the deviation is even more desirable for the same reason we articulated before. Now consider auxiliary problem (b) in addition. Since there is no updating before firm n switches back to the risky arm, taking $\tilde{T}_n = +\infty$ is also necessarily a profitable deviation. Therefore, the same condition in Step 1 will apply.