

Optimal Sharing Strategies in Dynamic Games of Research and Development¹

Nisvan Erkal² and Deborah Minehart³

First version: January 2005

This version: January 2006

The views expressed do not purport to represent the views of
the United States Department of Justice.

¹We have benefited from conversations with Eser Kandogan, Suzanne Scotchmer, Ben Shneiderman, and members of the Research Division of IBM. Nisvan Erkal thanks the Faculty of Economics and Commerce, University of Melbourne for its financial support.

²Department of Economics, University of Melbourne, Victoria 3010, Australia.

³United States Department of Justice, Washington, D.C. 20530, USA.

Abstract

This paper analyses the dynamic aspects of knowledge sharing in R&D rivalry. In a model where research projects consist of N sequential stages, our goal is to explore how the innovators' incentives to share intermediate research outcomes change with progress and with their relative positions in an R&D race. We consider an uncertain research process, where progress implies a decrease in the level of uncertainty that a firm faces. We assume that firms are informed about the progress of their rivals and make joint sharing decisions either before or after each success. Changes in the firms' absolute and relative positions affect their incentives to stay in the race and the expected duration of monopoly profits if they finish the race first. We show that firms always prefer to have sharing between their independent research units if they are allowed to collude in the product market. However, competing firms may have either decreasing or increasing incentives to share intermediate research outcomes throughout the race. If the lagging firm never drops out, the incentives to share always decrease over time as the research project nears completion. The incentives to share are higher earlier on because sharing has a smaller impact on each firm's chance of being a monopolist at the end of the race. If the lagging firm is expected to drop out, the incentives to share may increase over time. We also use our framework to analyze the impact of patent policy on the sharing incentives of firms and show that as patent policy gets stronger, sharing incentives may decrease or increase depending on whether or not the lagging firm has increased incentives to drop out.

1 Introduction

The ability to create knowledge-based assets plays an increasingly important role in determining firms' competitiveness in the market place. The goal of this paper is to analyze dynamic aspects of knowledge sharing in research and development (R&D) rivalry. Knowledge sharing is an important way in which firms can acquire the technological knowledge they need during their innovation process. Firms are likely to benefit from sharing knowledge with competitors. However, such alliances pose especially difficult challenges. This leaves us with the following question. When would we expect cooperation to emerge between competitors?

In the economics literature, knowledge sharing between rival firms has been the focus of many papers. These papers have mainly studied firms' incentives to share research outcomes at one point in time, either before the start of research, as in the case of research joint ventures, or after the development of a technology, as in the case of licensing.¹ In reality, the decision to share intermediate steps with rivals may be an integral part of a dynamic research process. Hence, the aim of this paper is to ask not whether but when firms prefer to share their research outcomes during a research process and what the emerging patterns of sharing activities are.

While sharing may cause researchers to benefit from each other's expertise and generate better ideas, it may also result in a reduction in the commercial value of their ideas. From a social welfare perspective, sharing of research outcomes is desirable because it results in less duplication. Hence, it is important to determine how close profit-driven firms come to maximizing welfare. In the economics literature, knowledge spillovers are stated as one of the most important reasons for rival firms to agree to share knowledge. However, from a dynamic perspective, another important aspect is uncertainty. The process of research is generally characterized by a high level of uncertainty in the beginning. Progress in research can be described as a decrease in the level of uncertainty that researchers face. Hence, one of the novel aspects of this project is to focus on the role uncertainty plays in the decisions to share knowledge and to analyze how firms' incentives to share research outcomes change during a

¹See, for example, Kamien (1989) on licensing, and Katz (1986), D'Aspremont and Jacquemin (1988) and Kamien, Muller and Zang (1992) on research joint ventures. Patenting and informal sharing between employees of firms are two other methods through which knowledge may be disseminated between firms. See, for example, Scotchmer and Green (1990) on early innovators' incentives to patent and Severinov (2001) on informal sharing between employees.

research process as the level of uncertainty they face decreases.

We assume that research projects consist of several sequential steps. Researchers cannot proceed to the next step before successfully completing the prior step. Moreover, they cannot earn any profits before completing all steps of the project. In a dynamic R&D process, firms' incentives to share change as their positions in the race change for two reasons. First, the expected duration of monopoly profits for the leading firm depends on the progress the firms make during the research process. Second, the probability that any two firms will be rivals in the product market changes with progress.

An important feature of the model is that we assume the different steps of research are symmetric in all respects except in regards to how far away they are from the end of the project. In other words, the options and technology available to the firms are the same in all steps of the research process. We deliberately assume that there are no spillovers during the research process. It has been stressed in the literature that firms may have higher spillover rates and bigger appropriability problems in earlier stages of research than in later stages of research.² Although the rate of spillovers may shape the incentives to share, we show that it is not the only relevant factor. Assuming that there are no spillovers between the research efforts of different firms allows us to focus on the role uncertainty plays in knowledge sharing.

We assume that firms are informed about the progress of their rivals and make joint sharing decisions either before or after each success. While sharing may cause researchers to benefit from each other's expertise and help them avoid wasteful duplication of R&D, it may also result in a reduction in the commercial value of their ideas. Because sharing decreases the lead of one firm, it reduces the expected profits that the leader derives from finishing the race first and being a monopolist for some period of time. This cost is even greater if, but for the sharing, the lagging firm would drop out of the race.

Hence, the decision to share and the pattern of sharing activities critically depend on the lagging firm's incentives to stay in the race in case of no sharing. The results reveal that firms always prefer to have sharing between their independent research units if they are allowed to collude in the product market. Under rivalry, the incentives to share intermediate research outcomes decreases monotonically with progress if the lagging firm is expected never to drop

²See, for example, Katz (1986), Katz and Ordover (1990), and Vonortas (1994).

out. The incentives to share are higher earlier on because there is more uncertainty earlier on. Sharing has a smaller impact on each firm's chance of being a monopolist at the end of the race.

If the lagging firm is expected to drop out, the incentives to share may increase with progress. This is because earlier in the research process the lagging firm may have a higher incentive to drop out and, hence, the leading firm may have a higher chance of eliminating rivalry by not sharing. We also illustrate that the incentives to share increase as the gap between the firms decreases.

We next use our framework to analyze the impact of patent policy on firms' sharing decisions. The strength of patent policy can have an important impact on firms' sharing decisions because it determines the costs of inventing around patented technologies. We show that if a strengthening in patent policy causes a change in the investment decision of the lagging firm at any of the asymmetric histories, then sharing incentives in general get weaker. Otherwise, they generally get stronger.

In addition to contributing to the literature on knowledge sharing, this paper is also related to the literature on the management of innovation (Aghion and Tirole, 1994a and 1994b). The design of an optimal R&D strategy is a multi-faced problem. Two aspects of this problem, regarding the intensity of the R&D effort and the riskiness of the R&D projects chosen by firms, have been dealt with extensively in the literature.³ In this paper we are interested in analyzing how the optimal strategies of firms change with progress. Other papers that have studied how firms' optimal strategies change over time in a dynamic model of R&D are Grossman and Shapiro (1986 and 1987), Cabral (2003) and Judd (2003). Grossman and Shapiro (1986 and 1987) analyze how firms vary their efforts over the course of a research project. In an infinite-period race, Cabral (2003) allow firms to choose between two research paths with different levels of riskiness. He shows that the leader chooses a safe technology and the laggard chooses a risky one. Judd (2003) shows that there is excessive risk-taking by innovators. Our paper differs from these papers because we analyze how firms' incentives to share and diversify change over the course of a research project.

³See Reinganum (1989) for a survey of the papers that focus on the intensity of firms' R&D efforts. Bhattacharya and Mookherjee (1986), Klette and de Meza (1986) and Dasgupta and Maskin (1987) analyze the riskiness of the research projects chosen by firms. Cardon and Sasaki (1998) analyze whether firms prefer to work on similar or different R&D paths.

The paper proceeds as follows. In the next section, we describe our set-up. In Section 3, we explore what happens if the firms are allowed to collude in the product market. In Section 4, we analyze the effect of competition on the dynamic sharing incentives of firms in a model with ex-post sharing contracts and 2 research steps. We consider the case of N research steps in Section 5. In Sections 6 and 7, we discuss extensions of our basic model with ex-ante sharing contracts and asymmetric firms respectively. After discussing the impact of patent policy on sharing incentives in Section 8, we conclude in Section 9.

2 Model

2.1 Research Environment

Since we are interested in the effect of competition on firms' incentives to share, we consider an environment with two firms, $i = 1, 2$, that each invest in a research project. On completion of the project, a firm can produce output in a product market. We assume the firms produce goods that may be either homogeneous or differentiated, and that they compete as duopolists in the product market.⁴

To capture the idea of progress, we assume that a research project has N distinct steps of equal difficulty. Hence, we assume that the firms divide the research project into different steps and that each firm defines the steps in the same way. A firm cannot start to work on the next step before completing the prior step, and all steps of the project need to be completed successfully before a firm can produce output. There is no difference between the steps in terms of the technology or the options available to the firms.

In the literature on multi-stage research, the phases of research are often thought of as qualitatively different. For example, there may be two steps identified as "research" and "development." We do not make this distinction, but rather we seek to derive endogenous differences in the research phases that result from dynamics in the decisions made by the firms. A basic intuition is that as firms approach the end of the research process, their decisions might increasingly reflect the impending rivalry.

We assume that each firm operates an independent research facility. We model research

⁴We assume that the firms conduct the research to solve the same technical problem. However, unmodelled differences in production technologies can still lead them to produce differentiated products.

activity using a Poisson discovery process. Time is continuous, and the firms share a common discount rate r . To conduct research, a firm must incur a flow cost c per unit of time.⁵ Investment provides a stochastic time of success that is exponentially distributed with hazard rate α . This implies that at each instant of time, the probability that the firm completes a step is α . After completing a step, a firm can immediately begin research on the next step. The successes of the two firms are statistically independent. To represent the progress made by the firms, we use the notation of a research history (k_1, k_2) , where k_i stands for the number of successes of firm i .

At each point in time prior to completing the project, a firm decides whether or not to invest. A decision not to invest is assumed to be irreversible and equivalent to dropping out of the game.⁶ Each firm is risk neutral and makes decisions to maximize its discounted expected continuation payoff given the strategy of the other firm. The payoff structures are more fully described below. A firm that drops out earns a continuation payoff of zero. Given the memorylessness nature of the Poisson process, if a firm is conducting research, it will not stop unless there is a change in its relative position in the research process. If the rival completes one of the steps successfully, the firm may decide to drop out of the game at this point. We implicitly assume that when one firm develops a step successfully, it does not result in any technological spillovers. The successful firm can either keep the innovation a secret or patent it. Patenting does not prevent the rival from developing a non-infringing technology that serves the same purpose.

We will also allow firms to share their research. If one firm has completed one (or more) steps that the other firm has not, the leading firm can share its research with the lagging firm. After sharing, both firms can proceed to the next research step.⁷ The timing of sharing decisions and the contracts that govern the sharing process are described below.

⁵We do not allow the firms to choose continuous levels of research effort in our basic model. This assumption can be motivated by presuming a fixed amount of effort that each firm can exert, which is determined by the capacity of its R&D division. As an example, Khanna and Iansiti (1997) explain that given the highly specialized nature of the R&D involved in designing state-of-the-art mainframe computers, firms in this industry find it very expensive to increase their number of researchers available to them. We relax this assumption later on and consider the case of continuous effort levels.

⁶Later, we may relax this assumption so that the decision not to invest can be reversed. We do not think that the qualitative nature of our results will change.

⁷Sharing in our model has the same impact as patenting in Scotchmer and Green (1990). In both models, the lagging firm can proceed to the next step after disclosure.

Regarding the information structure, we make the following assumptions. The firms will be able to share their research successes, but one firm cannot acquire the rival's innovation without such sharing. For example, a firm cannot observe the technical content of the rival's research without explicit sharing.⁸ Everything else in the game is common knowledge. In particular, firms observe whether their rival is conducting research as well as whether the rival has a success. Third parties such as courts also observe this information.

We next consider the product market competition and the sharing process before explaining how the firms' payoffs are represented.

2.2 Product Market Competition

After a firm completes all stages of the research process, it can participate in the product market. The firms produce goods that may either be homogeneous or differentiated to some degree by unmodelled differences in the production technologies. We represent the product market competition in the following reduced form way.

If both firms have completed the research project, they compete as duopolists and each earns a flow profit of $\pi^D \geq 0$ forever. If only one firm has completed the research project, the firm earns a monopoly flow profit of $\pi^M > 0$ as long as the other firm does not produce output. Here, $\pi^M > \pi^D$. As a benchmark, we will consider the case that the firms make production decisions to maximize their joint profits in the product market. This results in a joint flow profit of π^J where $\pi^J \geq 2\pi^D$ and $\pi^J \geq \pi^M$. We use the notation $\tilde{\pi}^D = \frac{\pi^D}{r}$, $\tilde{\pi}^M = \frac{\pi^M}{r}$, and $\tilde{\pi}^J = \frac{\pi^J}{r}$.

These payoffs are sufficiently flexible to capture various models of product competition. For example, if the firms produce homogeneous products and compete as Bertrand or Cournot competitors, then $\pi^J = \pi^M > 2\pi^D$. If the firms produce differentiated products, then $\pi^J > \pi^M$ and the relationship between π^M and $2\pi^D$ will depend on the degree of product differentiation that exists between the products. For low levels of product differentiation, $\pi^M > 2\pi^D$; for high levels of product differentiation, $\pi^M \leq 2\pi^D$.⁹

⁸Alternatively, we could assume that successful firms win immediate patents. A leading firm could then prevent a lagging firm from copying its research by enforcing its patent. If both firms complete the same step, they win non-infringing patents.

⁹The magnitudes of each of the profits π^D , π^J and π^M do not depend on the decisions taken during the research phase. In future research, we may relax this assumption.

As an example, consider a demand function of the type $q_i = (a(1 - \gamma) - p_i + \gamma p_j) / (1 - \gamma^2)$, where $0 < \gamma < 1$ so that the products are substitutes.¹⁰ The goods are more differentiated the higher is γ . It is possible to show that $\pi^M \leq 2\pi^D$ if and only if γ is sufficiently large.¹¹

From now on, we consider the case that there are $N = 2$ steps in the innovation process. In Section 5, we consider how our results extend to the case of an N -step innovation process.

2.3 Sharing of Research Outcomes

There are potential efficiencies in our model for firms to cooperate in the research stage. Suppose that one firm successfully completes a stage of research before the other firm does. We assume that the successful firm can costlessly share this knowledge with the other firm, thereby saving the lagging firm from having to continue to invest to complete the stage. From the point of view of social efficiency, such sharing will always be efficient because it prevents resources being spent to duplicate research results.

Because of the efficiencies of sharing, regulators in many countries encourage sharing arrangements, especially in the early stages of research. Firms may use a variety of contractual arrangements to govern the sharing process. There may be some legal restrictions, however, that prohibit sharing contracts that would inhibit competition in the output market. We want to consider firms incentives to share research using legal contracts. To this end, we want to classify contracts as either legal or illegal and limit our attention to legal contracts. However, even in our relatively simple dynamic framework, there are many contracts that might be written and it is not always obvious which ones we might want to classify as anti-competitive. Our first approach will be to consider two types of sharing contracts that are commonly observed in practice and have also been analyzed elsewhere in the literature. Later, we may consider a wider family of contracts.

The main of sharing contract we consider is *ex post sharing* or licensing, where the leading firm shares its results with the lagging firm in exchange for a fixed fee. Sharing will occur whenever the joint profits of the two firms are higher with sharing than without sharing. We do not place any restrictions on the fee, but we assume that the successful firm (the leader) makes a take-it-or-leave-it offer to the other firm (the follower). The leader, therefore, has all

¹⁰Singh and Vives (1984) show how these demand functions derive from particular consumer preferences.

¹¹The Hotelling models provide other examples of differentiated duopoly that can correspond to these profits.

of the bargaining power in the negotiation and will offer a fee that leaves the follower just indifferent between accepting and rejecting.¹²

Because the research project has 2 steps, there are six histories at which one firm has more knowledge than the other. These are the histories $(1, 0)$, $(0, 1)$, $(2, 0)$, $(0, 2)$, $(2, 1)$ and $(1, 2)$. We assume that if sharing occurs, it physically occurs instantly once one of these histories is reached. Given the memoryless nature of the Poisson process, this assumption is not very restrictive. We also assume that when a leading firm is more than one step ahead of the lagging firm, all the additional steps are shared, so that the lagging firm catches up to the leading firm. This is a simplifying assumption.

We also consider a second type of sharing contract, *ex ante sharing*. We assume that at the histories $(0, 0)$ and $(1, 1)$, the firms can make a joint decision about investing in the next research step and agree that once the step is completed, both firms will have access to the knowledge.¹³ The sharing agreement allows for contingent payments between the firms when the step is completed and the physical sharing of knowledge occurs. Knowledge sharing arrangements of this nature are often referred to as research joint ventures (RJV). Formally, we assume that the research technologies are not affected by the agreement. This means there are no synergies between the firms in the research process.¹⁴ Rather, the RJV is an agreement that allows both firms to have access to a success achieved by either one. Hence, it creates the opportunity for the firms to avoid wasteful duplication of R&D results. Alternatively, it allows the firms to agree to have one of the two facilities shut down altogether.¹⁵

¹²This division of bargaining power is appealing because it insures that each firm earns the full return to its research effort. Other divisions of bargaining power might also be considered. An existing literature considers how other divisions of bargaining power in licensing arrangements may affect the research incentives. See, for example, Katz and Shapiro (1985 and 1986), Shapiro (1985), Green and Scotchmer (1990 and 1995), Aghion and Tirole (1994), and D'Aspremont, Bhattacharya and Gerard-Varet (2000).

¹³In histories where the firms do not have the same number of successes, they can make a sharing agreement which involves both *ex post* and *ex ante* sharing. This is a common occurrence in RJV's, where the firms share their existing knowledge in an area in order to be able to work on the same research question together.

¹⁴Analyzing the collaborative R&D agreements of alliances and consortia registered under the National Cooperative Research Act in the US, Majewski (2004) shows that when the participants are direct competitors, they are likely to avoid spillovers.

¹⁵Such an asymmetric shut-down decision will never be optimal in this model, whether or not the shut-down decision is reversible.

2.4 Histories and Payoff Structures

To describe the game at any point in time, we need to specify how many firms are still active in the game, how many successes each active firm has, and whether there has been sharing. We use the following notation. Let (k_1, k_2) denote a research history where k_i is the number of steps that firm i has completed. The histories can be partially ordered so that (k_1, k_2) is *earlier* than (k'_1, k'_2) if and only if $k_i \leq k'_i$ for $i = 1, 2$, with strict inequality for at least one firm. In the following analysis, we refer to histories where $k_1 = k_2$ as symmetric histories and to those where $k_1 \neq k_2$ as asymmetric histories.

If a firm has dropped out of the game, we use X to denote this in the history. For example, to represent the history where firm 1 is working on the second step and firm 2 has dropped out of the game, we use $(1, X)$.¹⁶

Finally, to complete the description of the histories, we need to incorporate the sharing decisions into our notation for the history of the game. At symmetric histories, where there is no possibility of ex post sharing, we continue to use the notation (k, k) to denote that each firm has k successes. At asymmetric histories, we need to indicate whether the firms have made a sharing decision. At the instant that a firm achieves a success, we denote the history as (k_1, k_2) with $k_1 \neq k_2$. At this point, the firms make a sharing decision.¹⁷ If the firms share, the history becomes (k, k) where $k = \max\{k_1, k_2\}$. If the firms do not share, the history becomes (k_1, k_2, NS) . For example, consider the history $(2, 1)$. If the firms share, then the history becomes $(2, 2)$. If the firms do not share, then the history becomes $(2, 1, NS)$. In a continuation game at (k_1, k_2, NS) , the firms do not get another chance to share until the next innovation is achieved.

At any point during the research process, we denote the discounted expected continuation payoff of firm i starting at the history (k_1, k_2) by $V_i(k_1, k_2)$. This payoff is developed recursively from future continuation payoffs. Consider, for example, the continuation payoff of firm 1 at the history $(1, 0, NS)$, $V_1(1, 0, NS)$. Suppose there will not be any sharing between the firms

¹⁶We do not extend the partial ordering to histories where a firm has dropped out. This is because we will only need to refer to the ordering at histories where both firms are still active in the game.

¹⁷We assume that the sharing decision takes place in the same instant of time as the success, but we are using separate notation to capture the history before and after the sharing decision. The history (k_1, k_2) precedes the history (k_1, k_2, NS) in the partial ordering of histories. The two histories have the same ordering relative to all other histories in the game.

at any future history and the lagging firm will always choose to invest. If firm 1 develops the second step and firm 2 continues to invest after firm 1 develops the second step, then firm 1's continuation payoff is $V_1(2, 0, NS)$. If firm 2 develops the first step before firm 1 develops the second step and both firms stay in the game, then firm 1's continuation payoff is $V_1(1, 1)$. Hence, we have

$$\begin{aligned} V_1(1, 0, NS) &= \int_0^\infty e^{-(2\alpha+r)} (\alpha V_1(2, 0, NS) + \alpha V_1(1, 1) - c) dt \\ &= \frac{\alpha V_1(2, 0, NS) + \alpha V_1(1, 1) - c}{2\alpha + r} \end{aligned}$$

where the payoffs $V_1(2, 0, NS)$ and $V_1(1, 1)$ are similarly developed from future continuation payoffs.

After a firm has finished the research process, it earns continuation profits in the output market. To see how the payoffs are constructed, suppose firm 1 is the leading firm and firm 2 continues to research the second step. If there is no possibility of sharing, then we are at the history $(2, 1, NS)$. Firm 1 can produce output as a monopolist. In each instant of time, firm 2 has a probability α of success. As soon as firm 2 is successful, firm 1 starts to earn duopoly profits forever after. Hence, we have

$$\begin{aligned} V_1(2, 1, NS) &= \int_0^\infty e^{-(\alpha+r)} \left(\pi^M + \alpha \frac{\pi^D}{r} \right) dt \\ &= \frac{\pi^M + \alpha \frac{\pi^D}{r}}{\alpha + r}. \end{aligned}$$

If the firms decide to share, the continuation payoff of firm 1 is equal to $V_1(2, 2) = \frac{\pi^D}{r}$. The payoffs at other histories are developed similarly using recursion.

2.5 Example of Game Structure

To clarify the timing of decisions in the research phases of the game, consider the following illustration. As in Scotchmer and Green (1990), we use a discrete game tree as a stylized representation of the underlying continuous time model. We assume that firms share using the ex post licensing arrangements discussed above.¹⁸

¹⁸Scotchmer and Green (1990) include a rigorous justification for this representation based on a dominant strategy argument. That argument needs to be modified for our model in part because unlike them, we model the investment decision as irreversible. However, the basic idea is the same.

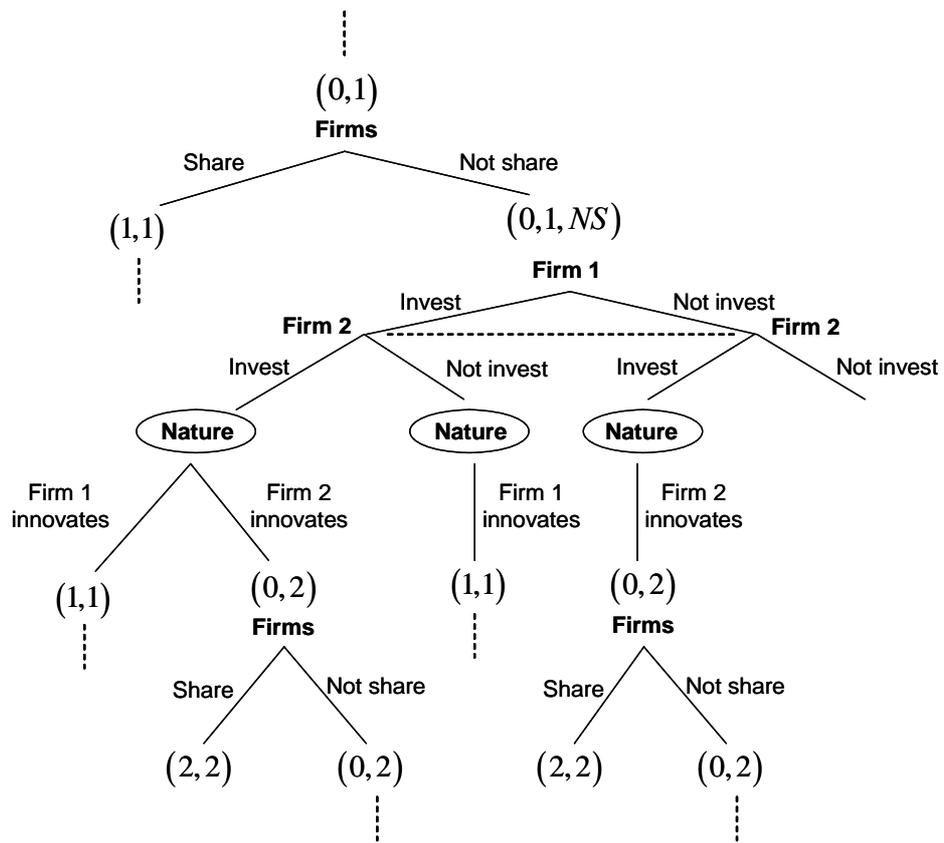


Figure 1: Game Tree following History $(0,1)$

At the beginning of the game, both firms simultaneously decide whether to invest or not. As shown in the subgame depicted in Figure 1, after one of the firms has made a discovery, the firms can jointly decide whether to share the winner's discovery. As mentioned before, sharing brings them to a symmetric position in the R&D phase. If they decide not to share, the laggard decides whether to continue to invest in order to develop the first innovation and the leader decides whether to invest to develop the second innovation. If they decide to share, they both simultaneously decide whether to invest in order to develop the second innovation.

If the laggard (re)develops the first innovation before the leader develops the second innovation, both firms start to invest to develop the second innovation. If the leader develops the second innovation before the laggard develops the first innovation, the firms again decide whether to share, this time both of the innovations. If they agree to share, they start to compete in the product market. If they agree not to share, the leader starts to produce while the laggard decides whether to continue to invest.

3 Benchmark Analysis

Before solving the game, we consider a benchmark case where the firms cooperate to maximize their joint profits. We assume that the firms make all investment, sharing, and product market decisions jointly.

At each history prior to the product market, the firms jointly decide whether to share the results of their research (if one firm is ahead) and whether one or both firms will invest.¹⁹ Once one firm completes the research project, the firms make joint decisions about the product market. We do not make any assumptions about how the firms divide the joint profits, but we simply assume that each decision is made to maximize the sum of the continuation profits of the two firms.²⁰ We have the following result.²¹

Proposition 1 *Suppose that the firms maximize their joint continuation profits. Then, at any history such that one firm has more research successes than the other, the optimal sharing*

¹⁹Throughout the paper, we assume that if a firm stops researching before completion of the project, it cannot reenter the game at a later date. This means that investment decisions are also participation decisions. We make the assumption for simplicity, but consider the consequences of relaxing it later.

²⁰We do not consider the possibility that the firms can commit to decisions prior to making them, but there is no dynamic inconsistency in the joint profit maximization problem.

²¹All the proofs are available upon request.

decision is for the leading firm to share its research with the lagging firm. Given this sharing pattern, the firms make investment decisions only at the symmetric histories $(0, 0)$ and $(1, 1)$. At the history $(2, 2)$, the firms cooperate in the product market and earn joint continuation profits $\tilde{\pi}^J = \frac{\pi^J}{r}$.

(i) If $\pi^J \geq \frac{2cr}{\alpha} + \frac{cr^2}{2\alpha^2}$, the optimal investment decision is for both firms to invest at the histories $(0, 0)$ and $(1, 1)$. At $(0, 0)$, the joint continuation profits are $\frac{4\alpha^2}{r(2\alpha+r)^2}\pi^J - \frac{2(4\alpha+r)}{(2\alpha+r)^2}c$.

(ii) If $\pi^J < \frac{2cr}{\alpha} + \frac{cr^2}{2\alpha^2}$, neither firm invests at $(0, 0)$ regardless of whether the firms would subsequently invest at $(1, 1)$. At $(0, 0)$, the joint continuation profits are 0.

Proposition 1 reveals several features of our model. First, it is jointly optimal for the firms to share a success as soon as it is developed by either one of them and to move on to the next stage of the R&D process. This reflects the traditional justification for sharing arrangements as a way for firms to avoid wasteful duplication. Second, if it is optimal for one firm to invest, it is optimal for both firms to invest. This is a feature of the Poisson discovery process that we are using. Indeed, with flow costs of investment, if there were N identical research facilities, then it would be optimal for all of them to conduct research simultaneously until one of the facilities achieves a success. This speeds up the time to innovation, and the benefits of the time savings outweigh the costs of running simultaneous facilities. Later, we discuss how our results might change if we used a model of the research process that does not have this feature.

The proposition illustrates how the cost and benefit parameters affect payoffs. In the region where firms invest, their joint continuation profits at the beginning of the game is increasing in α , the hazard rate for the Poisson discovery process. A higher α means that the research is likely to be successful sooner. The joint profits are also decreasing in the discount rate r and the flow cost of research c . Similar comparative statics results obtain for the continuation profits at other points in the game.

4 Ex-post Sharing

The first type of sharing contract we consider is *ex post sharing*. Suppose that one firm completes a research step ahead of the other firm. We consider a sharing contract where the leading firm shares its results with the lagging firm in exchange for a fixed fee. Such sharing takes place as long as it results in higher industry profits. Although a range of fees would

typically be acceptable to both firms, as discussed in Section 2.3, we will assume that the leader has all the bargaining power and sets the licensing fee by making a take-it-or-leave-it offer to the follower.

For each specification of the basic parameters of the game, we use backwards induction to find a subgame perfect equilibrium. We consider generic values of parameters, so that the equilibrium is unique and involves pure strategies. In this section, we first present the general properties of the equilibrium. Then, we present a full characterization for all possible values of π^M and π^D using a diagram and discuss how the equilibrium outcome changes as the profit levels change.

The primary benefit of sharing is that it avoids the wasteful duplication of R&D. The primary cost is the effect on output market competition. Because sharing erodes the lead of one firm, it reduces the expected profits that the leader derives from finishing the race first and being a monopolist for some period of time. This cost is even greater if, but for the sharing, the lagging firm would drop out of the race.²² The trade-off between monopoly and duopoly profits, has not been addressed at great length in the patent race literature because patent models usually assume a winner-take-all payoff structure.²³

A central question is how the incentives to share change over time. Because each of the research steps is identical from a technology standpoint, a conclusion that sharing incentives must change over time is not obvious. Certainly, if one firm is ahead of another, that may impact each firm's individual choices. However, if we consider the histories (1, 0) and (2, 1), it is not obvious that the sharing incentives should be any different. In both cases, the leader is one step ahead of the follower. Sharing is socially efficient in both cases and generates the same savings in terms of the elimination of wasteful R&D. The history (1, 0) is, however, earlier than the history (2, 1). At the earlier history, there is more uncertainty to be resolved before the firms enter the product market. We will consider how this uncertainty affects firms decisions.

The number of steps that the lagging firm is behind is also a factor in firms' sharing decisions. We control for this, however, by comparing histories such that the lagging firm is a

²²Scotchmer and Green (1990) examine the effect of secrecy on the drop-out decision of the firm.

²³Katz (1986) considers a model with one stage of research such that firms first engage in cooperative and independent R&D and then compete in an oligopolist output market. Also, see Cardon and Sasaki (1998) and Severnov (2001) for two more recent models of innovation where the firms compete as differentiated duopolists.

fixed number of steps behind the leading firm. This implies that in an N -step research process, we compare sharing incentives at all histories $(k + g, k)$ where g is a fixed gap between the leading firm and the lagging firm. The size of the gap can be as small as 1 or as large as $N - 1$. When $N = 2$, we compare the sharing decisions at $(2, 1)$ and $(1, 0)$ where the leading firm is one step ahead of the lagging firm.

From a dynamic perspective, what matters is whether one history precedes another. To consider dynamics, we compare histories $(k + g, k)$ and $(k' + g, k')$. If $k < k'$, then the history $(k + g, k)$ precedes the history $(k' + g, k')$. For example, when $N = 2$, the history $(1, 0)$ precedes the history $(2, 1)$.

The next definition states a formal monotonicity property for the general N -step model. When the property holds, sharing incentives may be said to decline over time as the firms approach the end of the game. We define the property for histories such that firm 1 is the leader. Because the equilibria in our game are symmetric, when the property holds, it also holds for histories such that firm 2 is the leader.

Definition 1 *An equilibrium satisfies the monotonicity property if whenever the firms share at the history $(k + g, k)$, then they also share at the earlier history $(k' + g, k')$ where $k' < k$. Here k and k' range from 0 to $N - g$ and $g = 1, \dots, N - 1$.*

We are now in a position to analyze the model when $N = 2$. Our central question is whether the monotonicity property holds. When $N = 2$, the monotonicity property states that whenever firms share at the history $(2, 1)$, they also share at the earlier history $(1, 0)$. There are three sharing patterns that satisfy this property. In the first pattern, the firms share at both $(1, 0)$ and $(2, 1)$. In the second pattern, the firms do not share at either $(1, 0)$ or $(2, 1)$. These patterns are weakly monotonic. In the third pattern, the firms share at $(1, 0)$, but they do not share at $(2, 1)$. This pattern is strongly monotonic. The monotonicity property fails if the firms do not share at $(1, 0)$, but do share at $(2, 1)$.

There are two principle motivations for firms to decide against sharing. First, if the lagging firm continues to research, it will take longer to finish the project allowing a longer expected period of monopoly profits for the leading firm. Second, if the lagging firm exits the game, the leader can expect to earn monopoly profits forever upon finishing. It turns out that

these two motivations can lead to different dynamics over time. To see this, we first consider environments such that exit does not occur.

Definition 2 *Region A consists of those parameter values such that in the unique subgame perfect equilibrium of the game, firms do not exit at any history either on the equilibrium path or off the equilibrium path. Region B consists of all other parameter values.*

Region A is given as follows:

Lemma 1 *Region A consists of all parameters such that $\pi^D \geq \frac{\sigma}{\alpha}(2 + \frac{r}{\alpha})$.*

Lemma 1 focuses on a firm that is as far behind the leader as possible when the leader has not shared its research. Because the lagging firm does not have any bargaining power, its payoff at $(2, 0, NS)$ or $(0, 2, NS)$ is the payoff it would get by conducting the two steps of research on its own and then producing in the output market as a duopolist. Intuitively, this is the worst possible position for a firm. We show that the lagging firm stays in at these histories if and only if the inequality $\pi^D \geq \frac{\sigma}{\alpha}(2 + r/\alpha)$ holds. We also show that when the inequality holds, the firms stay in the game at all other histories.

The next proposition records our main monotonicity result.

Proposition 2 *The monotonicity property holds for the unique subgame perfect equilibrium at all parameter values in Region A.*

The proof of the Proposition requires us to consider detailed equilibrium conditions. However, there is an underlying intuition for the result that we explain here. The benefit of sharing is the savings of duplicated R&D costs for one step of research. This benefit does not change over time. In contrast, the cost of sharing changes over time. The cost of sharing is measured in terms of the effect of sharing on future profits in the product market. Because of the resolution of uncertainty, sharing is more costly at $(2, 1)$ than it is at $(1, 0)$. To see this, note that in Region A, since the lagging firm never exits the game, the firms will be competing as duopolists in the product market eventually. At $(2, 1)$, the leading firm is done with the project and is assured of monopoly profits until the lagging firm catches up. If the firms share, then the leading firm foregoes these monopoly profits, as both firms immediately begin competing

as duopolists. In contrast, consider the earlier history $(1, 0)$. After sharing, the new history is $(1, 1)$. The leader has not foregone all of his chance to earn monopoly profits, because he can still finish the game first. In addition, at $(1, 0)$, the leader had no guarantee of monopoly profits anyway. The expectation of future monopoly profits is not so much lower after sharing than if the firms had not shared. Thus, sharing is less costly at $(1, 0)$ than at $(2, 1)$. In contrast, the benefits of sharing in terms of R&D cost savings do not change over time. The net effect is that firms are more likely to share earlier in the game. As we prove in Proposition 2, the incentive to share is always stronger at $(1, 0)$ than at $(2, 1)$. The intuition seems robust, so that we expect the monotonicity result to hold more generally for games with N -step research projects. We develop a result along these lines in Section 5.

We next consider region B. In this region, a lagging firm may exit the game if the leader does not share at some history. This introduces an important strategic motive for a leading firm to refuse to share. Our question is whether, in light of this, the pattern of sharing continues to satisfy the monotonicity property. We find that this is not the case. A lagging firm may be more likely to drop out earlier in the game, when it has more research left to complete. Given this, a leading firm may be less likely to share earlier in the game.

Proposition 3 *When $\frac{r}{\alpha} < 2$, there exist parameter values such that the subgame perfect equilibrium does not satisfy the monotonicity property. When $\frac{r}{\alpha} > 2$, the monotonicity property holds for the unique subgame perfect equilibrium at all parameter values in Region B.²⁴*

The monotonicity result of Proposition 2 extends to Region B, provided that $\frac{r}{\alpha} > 2$. A no sharing decision is never followed by a sharing decision. However, the monotonicity result cannot be further extended.

As we demonstrate in the proof of the proposition, when $\frac{r}{\alpha} < 2$, there exist values of π^D and π^M such that in equilibrium the firms do not share step 1 at $(1, 0)$, but do share step 2 at $(2, 1)$. The lagging firm drops out at the history $(1, 0, NS)$, but stays in at the later history $(2, 1, NS)$. The leader does not share at $(1, 0)$ because this would maintain a rival that is otherwise eliminated. Since the rival drops out at $(1, 0, NS)$, the history $(2, 1)$ is not reached

²⁴We prove the first statement in the Proposition, about non-monotonicity, in the appendix. The proof of the second statement in the Proposition, about monotonicity, is straightforward and available from the authors on request.

in equilibrium.²⁵ Hence, the non-monotonicity pattern is not observed along the equilibrium path. The parameter restriction means that the firms must be relatively patient and good at research for this type of equilibrium to exist. Sharing at $(1, 0)$ allows both firms to work on step 2 and hastens the end of the research phase. The firms are willing to forego this benefit when $\frac{r}{\alpha}$ is not too large.

We also demonstrate an equilibrium in which a non-monotonic sharing pattern arises on the equilibrium path. In equilibrium, the lagging firm drops out at the history $(2, 0, NS)$, but stays in at $(2, 1, NS)$ and $(1, 0, NS)$. The firms do not share step 1 at $(1, 0)$, because this way they can reach the history $(2, 0, NS)$, at which point the lagging firm drops out. A non-monotonic sharing pattern arises because after $(1, 0, NS)$, the firms sometimes reach the history $(1, 1)$. Both firms invest in step 2. The game may then proceed to the history $(2, 1)$, at which point the leading firm shares step 2 with the lagging firm.

For this type of equilibrium to exist, we need to impose a stronger restriction on $\frac{r}{\alpha}$. As the proof shows, for $\frac{r}{\alpha} > \frac{1}{2}(\sqrt{5} - 1)$, there are no non-monotonic sharing patterns along the equilibrium path. For $\frac{r}{\alpha} < \frac{1}{2}(\sqrt{5} - 1)$, there exist values of π^D and π^M such that the non-monotonic sharing pattern arises on the equilibrium path.

Figure 2 depicts the equilibrium outcomes in the case when we see non-monotonic sharing patterns both on and off the equilibrium path. This is the case that $\frac{r}{\alpha} < \frac{1}{2}(\sqrt{5} - 1)$. The diagram lists the sharing pattern for each region. For example, at the top left of the diagram, the sharing pattern S,NS,NS describes the following sequence of decisions: i) at $(1, 0)$, the leader shares (S) step 1; ii) at $(2, 0)$, the leader does not share (NS) step 1; iii) at $(2, 1)$, the leader does not share (NS) step 2. The diagram shows that there are two regions with the non-monotonic sharing pattern NS,NS,S. This pattern is non-monotonic because the leader does not share step 1 at $(1, 0)$, but does share step 2 at $(2, 1)$. The two regions with this sharing pattern are separated by a vertical line. In the region to the left of the line, the follower drops out at the history $(1, 0, NS)$. Because of this, the history $(2, 1)$ is not reached in the equilibrium.²⁶ Thus, an observer of the game would not observe a non-monotonicity. In the region to the right of the line, the follower stays in the game at the history $(1, 0, NS)$. Because of this, the history $(2, 1)$ is reached along the equilibrium path.

²⁵By symmetry, the follower also drops out at the history $(0, 1, NS)$. Thus, there is no path to $(2, 1)$.

²⁶By symmetry, the follower also drops out at the history $(0, 1, NS)$. Thus, there is no path to $(2, 1)$.

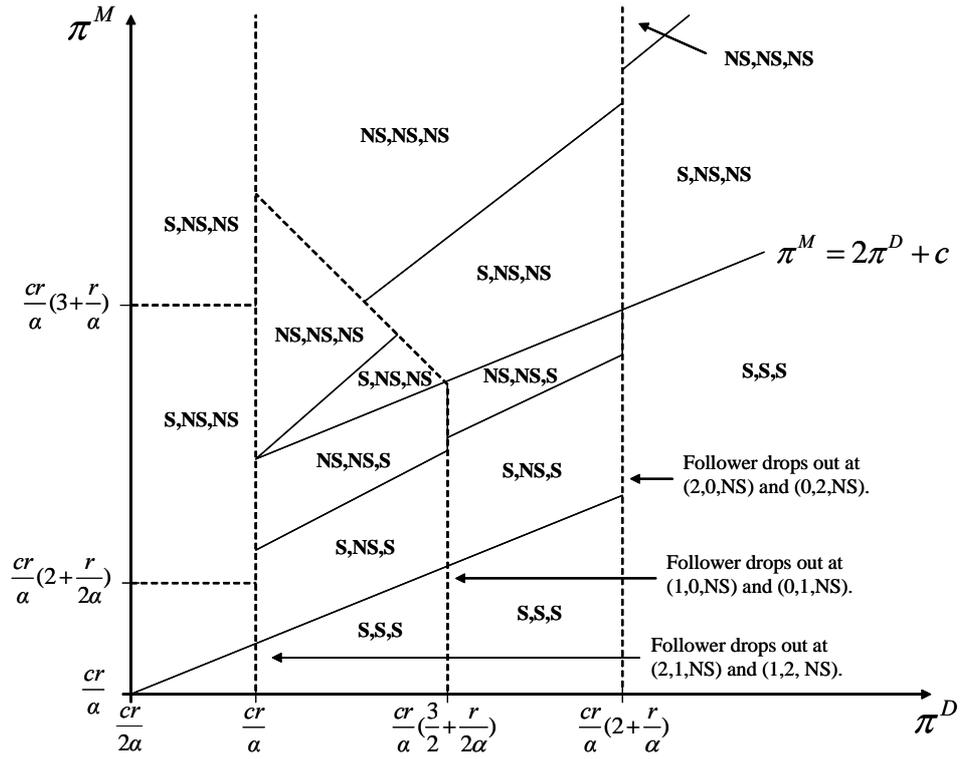


Figure 2: Equilibrium Outcomes

The diagram also shows the sharing patterns in other regions. Consider region A. Here, we have that $\pi^D \geq \frac{c\alpha}{\alpha}(2 + \frac{r}{\alpha})$. In Region A, the lagging firm never drops out of the game and the sharing patterns are monotonic. Of course, by Proposition 2, we already knew that this result must hold. Consider how the sharing pattern changes as π^M increases, but with the value of π^D held fixed. For small values of π^M , the sharing pattern is S,S,S. As monopoly profits increase, sharing breaks down at the history (2, 1).²⁷ As monopoly profits are increased further, sharing eventually breaks down at the earlier history (1, 0) as well. This is a monotonicity result for the comparative static analysis. As monopoly profits π^M increases, sharing breaks down, but it breaks down at later histories first.

Two regions in the diagram have a sharing pattern of S,NS,S. Here, the leading firm does not share step 1 at (2, 0), even though it does share step 1 at the earlier history (1, 0). We do not interpret this as a non-monotonicity result, because we only compare sharing decisions at histories where the gap between the leader and the follower is the same.

5 *N*-step Research Process

In this section, we discuss some results obtained in a model with N research steps of equal difficulty.

As a starting point, we consider our benchmark model where the firms cooperate to maximize their joint profits. The firms make all investment, sharing, and product market decisions jointly. Proposition 1 extends in a straightforward way. At asymmetric histories, the optimal sharing decision is for the leading firm to share its research with the lagging firm. At the final history (N, N) , the firms cooperate in the product market to earn joint continuation profits $\tilde{\pi}^J$. If the joint continuation profits are above a critical threshold, then both firms invest at all earlier symmetric histories. Otherwise, at the start of the game, neither firm invests.

We next turn to our monotonicity result that sharing declines over time. Proposition 2 extends to a model with 3 research steps.²⁸ We have

Proposition 4 *When $N = 3$, the monotonicity property holds for the unique subgame perfect*

²⁷Sharing also breaks down at the histories (2, 0) and (0, 2), but we do not have any other histories to compare these to with the same gap of 2 steps between the leader and the follower.

²⁸The proof is available from the authors on request. The calculations are straightforward, but long.

equilibrium at all parameter values in Region A. With $N = 3$, Region A is defined as the set of parameters such that $\pi^D \geq \frac{cr}{\alpha} (3 + 3\frac{r}{\alpha} + \frac{r^2}{\alpha^2})$.

To prove the proposition, we derive the equilibria as we did for the case of $N = 2$.²⁹ Region A is the set of parameters such that the lagging firm would stay in the game at the history $(3, 0, NS)$. At this history, the lagging firm is as far behind as possible and has no hope of ever earning monopoly profits. Because the lagging firm does not have any bargaining power, its payoff at $(3, 0, NS)$ is the payoff it would get by conducting all three steps of research on its own and then producing in the output market as a duopolist. Region A is all parameters such that this payoff is positive.

The monotonicity property implies that if the firms share at the histories $(k + 1, k)$, then they share at $(k, k - 1)$ for $k = 1, 2$. At these histories, the leading firm is one step ahead of the lagging firm. The monotonicity property also implies that at the history $(3, 1)$ when the leading firm is two steps ahead, they share at the earlier history $(2, 0)$.

The monotonicity property cannot be strengthened to comparisons between histories such that the leading firm is ahead by a differing number of steps. For example, we find an equilibrium such that the firms share at $(2, 1)$, but do not share at the earlier history $(2, 0)$. The reason is that at $(2, 0)$, the leading firm is further ahead and has more to give up in terms of forgone monopoly profits.³⁰

We expect that Proposition 2 could be extended further to a model with N research steps. The intuition behind the proposition is general and does not depend on the assumption that $N = 2$. The benefit of sharing is the savings on R&D costs. These cost savings do not change over time. However, the cost of sharing, measured in terms of foregone monopoly profits, do change over time. The advantage to the leading firm to being a fixed number of steps ahead of the lagging firm increases over time as uncertainty is resolved. The net effect is that the firms have decreasing incentives to share as the game progresses

We have not proved Proposition 2 for the general case of N research steps, because the equilibrium calculations become too cumbersome. Instead, we analyzed a related problem

²⁹The proof is available from the authors on request.

³⁰When $N = 2$, we also found equilibria in Region B such that the firms share at $(2, 1)$, but not at $(2, 0)$. The sharing pattern arises because the lagging firm exits after the decision not to share. When $N = 3$, we find this sharing pattern in Region A, even though the lagging firm does not ever exit the game.

that we interpret as a partial generalization of our monotonicity result. Consider any starting history $(k + 1, k)$ in the N -step model such that the leading firm is one step ahead of the lagging firm. If the firms share at this history, then the new history becomes $(k + 1, k + 1)$. If they do not share, then the new history is $(k + 1, k, NS)$. Assume that at all histories after $(k + 1, k, NS)$ and $(k + 1, k + 1)$ the firms do not share and they also do not exit the game. Under this assumption, we can derive formulas representing the firms' joint continuation payoffs. We can compare the continuation payoffs from sharing and not sharing at $(k + 1, k)$. The benefit of sharing (which is equal to the cost savings by the lagging firm) is an increasing (linear) function of the flow cost of research c . Because of this, there is a threshold cost $c(k + 1, k)$ such that the firms decide to share if and only if $c \geq c(k + 1, k)$. Using numerical analysis, we can show that the threshold cost $c(k + 1, k)$ is increasing in k for $N \leq 20$.³¹

The finding suggests that sharing is more likely to occur at earlier histories. The cost parameter c is more likely to be above the sharing threshold cost $c(k + 1, k)$ at earlier histories than at later ones. The result is different from Proposition 2 because the assumptions about firms' behavior after $(k + 1, k)$ may not be consistent with any equilibrium. However, the result is consistent with the intuition that the incentives to share decline over time when firms never exit.

Our other result, Proposition 3, is that there are parameter values in Region B such that a subgame perfect equilibrium does not satisfy the monotonicity property. These non-monotonic equilibria continue to exist as subgames of the N -step model. This is because the subgame that begins at the history $(N - 2, N - 2)$ is a two-stage game. Parameters that support the equilibrium are in Region B of the 2 step game. The parameters are also in Region B of the N -step game because Region B grows as N increases.³² Thus, Proposition 3 continues to hold.³³

³¹We compared the payoffs by evaluating them on a discrete grid of parameter values. The formulas appear to be sufficiently continuous that we do not expect we missed any singularities in our simulations. The computations are available on request.

³²Region A shrinks as N increases, because a lagging firm has a lower payoff from staying in the game at $(N, 0)$ than at $(N - 1, 0)$. Region B grows as Region A shrinks, since they are complementary sets.

³³If a firm drops out of the N -stage game prior to the history $(N - 2, N - 2)$, then the continuation equilibrium would still exist but would represent off-the-equilibrium path behavior.

6 Ex-ante Sharing

In this section, we consider a second type of sharing contract. We assume that at any (symmetric) history, the firms can make a joint decision about investing in the next research step and agree that once the step is completed, both firms will have access to the knowledge as in a RJV.

Consider a sharing contract that is signed at the beginning of the game. The history is $(0, 0)$. The firms both agree to conduct research on the first step. The firms also agree that when one firm has a success, it will share the success with the lagging firm. In exchange, the lagging firm will pay a fee to the leading firm. At the time the contract is signed, the fee is contingent - it is paid by the lagging firm to the leading firm at the instant of innovation. We assume that the fee is set so that the lagging firm is indifferent between paying the fee to get the result, and not paying the fee and not getting the result. This means that the leading firm extracts the full value of its success. Therefore, as was the case under our ex post sharing contracts, both firms have efficient incentives to invest ex ante.³⁴

We consider a similar sharing contract at the history $(1, 1)$. We then analyze the sharing pattern to find whether the firms are more likely to share at $(0, 0)$ than at $(1, 1)$. Our analysis shows that there is no difference in terms of sharing incentives between the ex ante sharing contracts and the ex post sharing contracts. Because of this, the dynamics of sharing over time are essentially the same as before.

7 Asymmetric Firms

So far we have assumed that the firms are symmetric in their research capabilities. In this section we relax this assumption and assume that the firms can differ in their research capabilities. That is, different firms may have different areas of expertise which make them perform better in different stages of the research process. This is often the case, for example, in the biotechnology industry (Greis et al., 1995). Large pharmaceutical corporations form alliances with small research firms which perform the basic research towards the development of a new product. After small research firms successfully complete the initial stages of research project,

³⁴The contract need not provide any direct incentives for investment and would induce efficient effort even if the effort levels were non-contractible.

large pharmaceutical corporations work to bring the new product into the market.

To represent this kind of a scenario, we consider the $N = 2$ model and assume that one of the firms in our model is better at first-stage research and the other firm is better at second-stage research. Let α stand for the hazard rate of the more capable firm and β stand for the hazard rate of the less capable firm in each period.

Recall from our benchmark analysis in Section 3 that if the firms can collude in the output market, then the optimal sharing decision is for the leading firm to share its research with the lagging firm. Moreover, if it is optimal for one firm to invest, it is optimal for both firms to invest. With asymmetric firms, it is still the case that the optimal sharing decision is for the leading firm to share its research with the lagging firm. However, as far as the investment incentives of the firms are concerned, we can have one of the following outcomes: (i) Neither firm invests, (ii) both firms invest in both stages, (iii) only the efficient firm invests in both stages, and (iv) only the efficient firm invests in the first stage and both firms invest in the second stage. There are no equilibrium outcomes where only the inefficient firm invests because if expected profits are high enough for the less efficient firm to invest, they are high enough for the more efficient firm to invest.

Under rivalry, the monotonicity result stated in Proposition 2 for Region A still holds. The boundary of Region A is again defined by the incentives for the lagging firm to stay in the race at the histories $(2, 0, NS)$ or $(0, 2, NS)$. We get the same condition whether it is firm 1 or firm 2 that is the lagging firm.

In Region B, the asymmetry between the firms causes an asymmetry in the firms' drop-out decisions. That is, there may be cases where the more efficient firm, if it were the lagging firm, would stay in the market, but the less efficient would not.

As far as the sharing conditions are concerned, since the hazard rate does not enter the second-stage sharing condition, asymmetry does not affect the sharing decision at the histories $(2, 1)$, $(1, 2)$, $(2, 0)$ and $(0, 2)$. After the first success, there may be cases where there would be sharing if it is the more efficient firm that has the first success, but no sharing if it is the less efficient firm that has the first success. If it is the more efficient firm that has the first success, this implies that the firm that is less efficient in the second-stage research is ahead and the

firm that is less efficient in first-stage research is behind. Both strengthen the incentives to share the first success.

With the consideration of asymmetric firms, one question that arises is whether we may have cases where only the efficient firm invests in each stage. To consider this issue we need to change the assumption we have made so far that once a firm exits the race, it cannot re-enter. Allowing re-entry, we can show that firms with asymmetric hazard rates may find it optimal to specialize. To see this, consider the extreme case where $\beta = 0$. It is straightforward to show that only the efficient firm invests at $(0, 0)$. As soon as the first success arrives, the firm shares it with the non-investing firm and drops out. Hence, without the prospective of sharing the first success at $(1, 0)$, the firm would not have invested at $(0, 0)$.

8 Impact of Patent Policy

The framework we have developed can be used to investigate several policy questions. Two important questions in patent policy are how strong patent protection should be and how strict the non-obviousness requirement should be (i.e., whether intermediate research outcomes should be patentable). In this section, we focus on the first question.

So far we have assumed that once a firm successfully develops a research step, it can either keep the technology a secret or patent it. Patenting does not prevent the rival from developing a non-infringing technology that serves the same purpose. In reality, the impact of patenting on the rival's progress would depend on the strength of patent policy. Stronger patent policy would make it harder for rival firms to invent around (Gallini, 1992).

This implies that the strength of patent policy can have an important impact on firms' sharing decisions. Hence, we next use our framework to analyze the impact of patent policy on firms' sharing decisions. Consider a variation of our basic model where firms can choose between a continuum of research paths. Different research paths are associated with different hazard rates. Firms still must incur a flow cost c per unit of time if they decide to invest. If firm i decides to invest, it can choose a research path that yields a hazard rate $\alpha_i \in [0, \bar{\alpha}]$. We assume that the research paths can be ranked in terms of their quality (i.e., how promising they are) and that both firms rank the steps in the same way. Hence, $\bar{\alpha}$ represents the hazard rate that is associated with the most promising research path.

Clearly, if all of the research paths are available, the firms would choose the research path that yields the highest hazard rate, $\bar{\alpha}$, at the beginning of the race. After one of the firms is successful, it patents the new technology. Patenting implies that if the rival continues to invest, it must switch to a different research path, where it faces a lower hazard rate.³⁵ Let α^L denote the maximum hazard rate that the lagging firm can achieve without infringing the patent of the leader. A stronger patent policy can be interpreted as corresponding to a lower α^L . Hence, we are interested in investigating how the sharing incentives change as α^L decreases.

We divide the analysis into two parts depending on whether a strengthening in patent policy changes the investment decision of the lagging firm at any of the asymmetric histories. If patent policy gets stronger without affecting the dropping out decision of the lagging firm at any of the asymmetric histories, it can have two kind of effects on the sharing incentives. First, since a decrease in α^L makes the lagging firm less efficient, it increases the benefits from sharing. Second, since a decrease in α^L strengthens the leader's position, it increases the costs of sharing. The first effect dominates everywhere in Region B. It also dominates in Region A for sufficiently small changes in patent policy. For larger changes in patent policy (i.e., larger decreases in α^L), the second effect dominates because the rival becomes considerably weaker.

If a strengthening in patent policy increases the set of histories where the lagging firm drops out, then sharing incentives in general get weaker. This is because as patent policy gets stronger, the position of the lagging firm gets weaker and it is more likely to drop out. This reduces the sharing incentives.

However, this reasoning does not apply in the case when a strengthening in patent policy causes the lagging firm to drop out at $(2, 1)$ and $(1, 2)$. In this case, the incentives to share step 1 at the histories $(1, 0)$, $(0, 1)$, $(2, 0)$ or $(0, 2)$ may get stronger relative to the case when the lagging firm would not drop out at $(2, 1)$ and $(1, 2)$.³⁶ This is because if the lagging firm will not drop out at $(2, 1)$ and $(1, 2)$, the leader chooses not to share the first step in order to prolong the time period during which it can potentially earn monopoly profits. If the lagging firm will drop out at $(2, 1)$ and $(1, 2)$, the firm that first successfully develops the first step does not have to protect itself by not sharing.

³⁵Note that due to the memoryless nature of the Poisson discovery process, such switching can happen independent of the lagging firm's initial choice of research path.

³⁶More specifically, this is the case for sufficiently high values of π^M .

9 Conclusion

The paper considers the optimal pattern of knowledge sharing in the context of technological competition. Developing a theoretical foundation for optimal sharing strategies has important implications for the design of optimal as well as efficient research environments.

We have analyzed how the incentives to share change over time as a research project reaches maturity. The decision to share and the pattern of sharing activities critically depend on the lagging firm's incentives to stay in the race in case of no sharing. The results reveal that under rivalry, the incentives to share intermediate research outcomes decrease monotonically with progress if the lagging firm is expected never to drop out. The incentives to share are higher earlier on because there is more uncertainty earlier on. Sharing has a smaller impact on each firm's chance of being a monopolist at the end of the race.

In many models of R&D, there is an assumption that firms share at an early research stage but not at a later one. This result shows that this sharing pattern can be derived from the optimizing behavior of firms in a dynamic game where the research technology does not change over time.

If the lagging firm is expected to drop out, the incentives to share may increase with progress. This is because earlier in the research process the lagging firm may have a higher incentive to drop out and, hence, the leading firm may have a higher chance of eliminating rivalry by not sharing.

An analysis of the impact of patent policy on firms' sharing decisions reveals that sharing incentives in general get weaker if a strengthening in patent policy causes a change in the investment decision of the lagging firm at any of the asymmetric histories. Otherwise, sharing incentives generally get stronger. The framework can also be used to analyze whether patentability of intermediate research outcomes is desirable and under what circumstances a more tolerant treatment of research collaborations may be desirable.

References

- [1] Aghion, P. and J. Tirole. 1994a. "The Management of Innovation," *Quarterly Journal of Economics*, 109(4), 1185-1209.
- [2] Aghion, P. and J. Tirole. 1994b. "Opening the Black Box of Innovation," *European Economic Review*, 38, 701-710.
- [3] D'Aspremont, C., S. Bhattacharya and L.-A. Gerard-Varet. 2000. "Bargaining and Sharing Innovative Knowledge," *Review of Economic Studies*, 67, 255-271.
- [4] D'Aspremont, C. and A. Jacquemin. 1986. "Cooperative and Noncooperative R&D in Duopoly with Spillovers," *American Economic Review*, 78(5), 1133-1137.
- [5] Bhattacharya, S. and D. Mookherjee. 1986. "Portfolio Choice in Research and Development," *Rand Journal of Economics*, 17(4), 594-605.
- [6] Cabral, L. M. B. 2003. "R&D Competition When Firms Choose Variance," *Journal of Economics and Management Strategy*, 12(1), 139-150.
- [7] Cardon, J. H. and D. Sasaki. 1998. "Preemptive Search and R&D Clustering," *Rand Journal of Economics*, 29(2), 324-338.
- [8] Dasgupta, P. and E. Maskin. 1987. "The Simple Economics of Research Portfolios," *The Economic Journal*, 97, 581-595.
- [9] Green, J. and S. Scotchmer. 1995. "On the Division of Profit in Sequential Innovation," *Rand Journal of Economics*, 26, 20-23.
- [10] Grossman, G. M. and C. Shapiro. 1986. "Optimal Dynamic R&D Programs," *Rand Journal of Economics*, 17(4), 581-593.
- [11] Grossman, G. M. and C. Shapiro. 1987. "Dynamic R&D Competition," *Economic Journal*, 97(386), 372-387.
- [12] Judd, K. L. 2003. "Closed-Loop Equilibrium in a Multi-Stage Innovation Race," *Economic Theory*, 21, 673-695.

- [13] Kamien, M. I. 1992. "Patent Licensing," in R. J. Aumann and S. Hart (eds.), *Handbook of Game Theory with Economic Applications*, Volume 1, Elsevier Science.
- [14] Kamien, M. I., E. Muller and I. Zang. 1992. "Research Joint Ventures and R&D Cartels," *American Economic Review*, 82(5), 1293-1306.
- [15] Katz, M. L. 1986. "An Analysis of Cooperative Research and Development," *Rand Journal of Economics*, 17(4), 528-542.
- [16] Katz, M. and C. Shapiro. 1985. "On the Licensing of Innovations," *Rand Journal of Economics*, 16, 504-520.
- [17] Katz, M. and C. Shapiro. 1986. "How to License Intangible Property," *Quarterly Journal of Economics*, 101, 567-589.
- [18] Khanna, T. 1998. "The Scope of Alliances," *Organization Science*, 9(3), 340-355.
- [19] Khanna, T. and M. Iansiti. 1997. "Firm Asymmetries and Sequential R&D: Theory and Evidence from the Mainframe Computer Industry," *Management Science*, 43(4), 405-421.
- [20] Klette, T. and D. de Meza. 1986. "Is the Market Biased Against Risky R&D?" *Rand Journal of Economics*, 17(1), 133-139.
- [21] Lee, T. and L. L. Wilde. 1980. "Market Structure and Innovation: A Reformulation," *Quarterly Journal of Economics*, 94(2), 429-436.
- [22] Majewski, S. 2004. "How Do Consortia Organize Collaborative R&D? Evidence from the National Cooperative Research Act," Harvard Law School, Olin Center for Law, Economics, and Business, Discussion Paper No. 483.
- [23] Oxley, J. and R. Sampson. 2004. "The Scope and Governance of Knowledge-Sharing Alliances," *Strategic Management Journal*, 8-9, 723-750.
- [24] Reinganum, J. 1989. "The Timing of Innovation," in R. J. Aumann and S. Hart (eds.), *Handbook of Industrial Organization*, Volume 1, Elsevier Science.
- [25] Scotchmer, S. and J. Green. 1990. "Novelty and Disclosure in Patent Law," *Rand Journal of Economics*, 21(1), 131-146.

- [26] Severinov, S. 2001. "On Information Sharing and Incentives in R&D," *Rand Journal of Economics*, 32(3), 542-564.
- [27] Shapiro, C. 1985. "Patent Licensing and R&D Rivalry," *American Economic Review*, 75(2), 25-30.

Appendix

A Proof of Proposition 1

We solve the model under the assumption that the two firms maximize their joint payoffs. We derive continuation profits at each history working backwards through the decision nodes.

At (2, 2), the firms cooperate in the output market to earn the joint flow profit π^J forever. Recall that $\pi^J \geq \max\{2\pi^D, \pi^M\}$ so that profits in the output market are greatest when the firms produce cooperatively. The joint continuation profits are:

$$V_J(2, 2) = \frac{\pi^J}{r} = \tilde{\pi}^J$$

At the histories³⁷ (2, 1) and (2, 0), the leading firm shares all available research with the lagging firm. This prevents the wasteful duplication of R&D. The firms then cooperate in the product market to earn joint continuation profits of $\tilde{\pi}^J$. Thus, we have that

$$V_J(2, 1) = V_J(2, 0) = V_J(2, 2) = \tilde{\pi}^J$$

At the history (1, 1), if neither firm invests, the joint continuation profits are 0. If one firm invests (either firm), then the firm invests a flow cost of c and in each instant the probability of success is α . When the success arrives, the firms share the research and cooperate in the product market to earn flow profits of π^J . At (1, 1), the joint continuation profits are:

$$V_J(1, 1) = \int_0^\infty e^{-(\alpha+r)t} (\alpha\tilde{\pi}^J - c) dt = \frac{\alpha\tilde{\pi}^J - c}{\alpha + r}$$

If both firms invest, then each firm incurs a flow cost of c and the flow probability that at least one firm succeeds is 2α . The joint continuation profits are:

$$V_J(1, 1) = \int_0^\infty e^{-(2\alpha+r)t} (2\alpha\tilde{\pi}^J - 2c) dt = \frac{2\alpha\tilde{\pi}^J - 2c}{2\alpha + r} \quad (1)$$

Given these payoffs, the firms will either both invest or both not invest. The firms invest³⁸ if and only if $V^J(1, 1) \geq 0$. This occurs if and only if $\tilde{\pi}^J \geq \frac{c}{\alpha}$.

³⁷The histories (1, 2) and (0, 2) are analyzed in the same way as (2, 1) and (2, 0). We do not repeat the analysis here.

³⁸For simplicity (but with some abuse of notation), we ignore (non-generic) parameters such that the firms are indifferent between two actions.

Working backwards, we reach the history (1, 0). As at (2, 0) and (2, 1), sharing eliminates wasteful duplication of R&D. Because the firms make decisions cooperatively, there is no cost to them to sharing. Sharing either strictly increases their joint continuation profits or has no effect on the profits because the firms are in any event exiting the race. Without loss of generality, we will assume that the firms share at (1, 0).

Finally, we consider the history (0, 0). If neither firm invests, their joint continuation profits are 0. If both firms invest, then their joint continuation profits are:

$$V_J(0, 0) = \int_0^\infty e^{-(2\alpha+r)}(2\alpha V_J(1, 1) - 2c)dt = \frac{2\alpha V_J(1, 1) - 2c}{2\alpha + r}$$

The continuation profits depend on whether the firms invest at (1, 1). If the firms do not invest at (1, 1), then they clearly will not invest at (0, 0). If the firms invest at (1, 1), then they invest at (0, 0) if and only if $V_J(0, 0) \geq 0$. This is the case if and only if $V_J(1, 1) \geq \frac{c}{\alpha}$. Using the expression for $V_J(1, 1)$ above, we find that the firms invest at (0, 0) if and only if $\pi^J \geq \frac{cr}{\alpha}$ and

$$\pi^J \geq \frac{2cr}{\alpha} + \frac{cr^2}{2\alpha^2}.$$

The last inequality above implies the inequality $\pi^J \geq \frac{cr}{\alpha}$. Thus, this inequality is a necessary and sufficient condition for both firms to invest at (0, 0). If the firms invest, then their joint continuation profits are

$$V_J(0, 0) = \frac{2\alpha V_J(1, 1) - 2c}{2\alpha + r} = \frac{4\alpha^2}{r(2\alpha + r)^2} \pi^J - \frac{2(4\alpha + r)}{(2\alpha + r)^2} c.$$

B Proof of Lemma 1

In a companion Appendix that is available on request, we analyze all the equilibria of the game. That analysis also proves the lemma. Here we take a different approach. We focus on the payoff for the lagging firm at the history (2, 0, NS). Intuitively, the history (2, 0, NS) represents the worst possible position for the lagging firm. At this history, the lagging firm is as far behind as possible and has no hope of ever earning monopoly profits. We argue that if the computed payoff is positive, the firm stays in the game at all histories and we are in Region A. Conversely, if the payoff is negative, the firm exits at (2, 0, NS) and we are not in Region A. We further show that the payoff is the same whether or not the firms share at (2, 1). Because of this, Region A is defined by the single condition given in Lemma 1.

We first argue that the payoff of the lagging firm at $(2, 0, NS)$ is a lower bound on any firm's payoff at any history and in any equilibrium. Because the lagging firm has no bargaining power, its payoff at $(2, 0, NS)$ is exactly the same as what it would earn by conducting two steps of research on its own and then producing in the output market as a duopolist. This payoff must be a lower bound on any firm's payoff at any history, because a firm always has the option of conducting research on its own and then producing in the output market to earn at least duopoly profits.

We next compute the payoff $V_2(2, 0, NS)$ to establish the condition that defines Region A. We work backwards through the game.

At $(2, 2)$, each firm produces output and makes $V_2(2, 2) = \tilde{\pi}^D = \frac{\pi^D}{r}$.

At $(2, 1, NS)$, firm 1 produces output. Firm 2 invests if and only if

$$V_2(2, 1, NS) = \frac{\alpha V_2(2, 2) - c}{\alpha + r} = \frac{\alpha \tilde{\pi}^D - c}{\alpha + r} > 0. \quad (2)$$

Therefore, firm 2 invests at $(2, 1, NS)$ if and only if $\tilde{\pi}^D > \frac{c}{\alpha}$. If firm 2 does not invest then its payoff is 0.

At $(2, 1)$, the leading firm retains all of the surplus from any sharing agreement, by assumption. Since the lagging firm has no bargaining power, it receives the same payoff whether or not it shares. Therefore $V_2(2, 1) = V_2(2, 1, NS)$.

At $(2, 0, NS)$, if the lagging firm invests, its continuation profit is:

$$V_2(2, 0, NS) = \frac{\alpha V_2(2, 1) - c}{\alpha + r} = \frac{\alpha V_2(2, 1, NS) - c}{\alpha + r}.$$

If $\tilde{\pi}^D < \frac{c}{\alpha}$, the lagging firm exits the race at $(2, 1, NS)$ and earns a profit of 0. Given this, the firm exits the race at $(2, 0, NS)$ as well. If $\tilde{\pi}^D > \frac{c}{\alpha}$, the lagging firm stays in the race at $(2, 1, NS)$. Substituting for $V_2(2, 1, NS)$ using (2), we have:

$$V_2(2, 0, NS) = \frac{\alpha \left(\frac{\alpha \tilde{\pi}^D - c}{\alpha + r} \right) - c}{\alpha + r}.$$

This payoff is strictly positive if and only if

$$\pi^D > \frac{cr}{\alpha}(2 + r/\alpha).$$

This is the inequality that defines Region A. To see this, first suppose that $\pi^D < \frac{cr}{\alpha}(2+r/\alpha)$. If the lagging firm exits the game at $(2, 1, NS)$, then we are not in Region A. If the lagging firm stays in the game at $(2, 1, NS)$, then its payoff if it invests at $(2, 0, NS)$ is the one we just computed. Since this is negative, the firm exits the game and we are not in Region A.

Next, suppose that $\pi^D > \frac{cr}{\alpha}(2+r/\alpha)$. At $(2, 1, NS)$, the lagging firm stays in the game because $\pi^D > \frac{cr}{\alpha}$. At $(2, 0, NS)$, the lagging firm stays in the game because $\pi^D > \frac{cr}{\alpha}(2+r/\alpha)$. Thus the lagging firm stays in the game at $(2, 0, NS)$, and we are in Region A. This is what we wanted to show.

C Proof of Proposition 2 and Proposition 3

We solve for the equilibria of the game for all parameter values in the companion Appendix to this paper. That analysis proves propositions 2 and 3. Figure 2 illustrates the equilibria for all values of π^D and π^M when $\frac{r}{\alpha} < \frac{1}{2}(\sqrt{5}-1)$. For readers who do not wish to read the companion Appendix, we show how to derive one of the equilibria below. The equilibria for other parameter values are solved similarly.

D Derivation of a Non-Monotonic Equilibrium

We solve the game in the following region of parameters: $\frac{cr}{\alpha}(\frac{3}{2} + \frac{r}{2\alpha}) < \pi^D < \frac{cr}{\alpha}(2 + \frac{r}{\alpha})$ and $2\pi^D(\frac{2\alpha+2r}{2\alpha+r}) + c(\frac{2r}{2\alpha+r}) < \pi^M < 2\pi^D + c$. This is a subregion of region B. A straightforward calculation³⁹ shows that the region is non-empty if and only if $\frac{r}{\alpha} < \frac{1}{2}(\sqrt{5}-1)$ where $\frac{1}{2}(\sqrt{5}-1) \simeq 0.62$. The equilibrium is also derived in the companion Appendix to this paper, where the region is labeled Region 6.

To find an equilibrium, we work backwards from the end of the game. We derive the continuation profits at each history and solve for the equilibrium actions. At asymmetric histories such as $(2, 1)$ and $(1, 2)$, the analysis of the game is the same so we analyze only one of the histories.

The last history is the history $(2, 2)$. At this history, the firms have two successes each and are done with the research. **They produce output and each earns discounted duopoly profits of $V_i(2, 2) = \tilde{\pi}^D = \frac{\pi^D}{r}$.**

³⁹This derivation is available on request.

Working backwards, **the next history is** $(2, 1, NS)$. **Firm 1 is finished with its research and produces output.** Firm 2 has 1 success, and the firms have declined to share. Firm 2 decides whether or not to invest in step 2. If firm 2 invests, then its continuation profit is

$$V_2(2, 1, NS) = \int_0^\infty e^{-(\alpha+r)}(\alpha V_2(2, 2) - c)dt = \frac{\alpha\tilde{\pi}^D - c}{\alpha + r}. \quad (3)$$

This payoff is positive because by assumption

$$\pi^D > \frac{cr}{\alpha}$$

Hence **firm 2 invests at** $(2, 1, NS)$. Firm 1 earns monopoly profits until firm 2 completes the second step. The continuation profit of firm 1 is

$$V_1(2, 1, NS) = \int_0^\infty e^{-(\alpha+r)}(\pi^M + \alpha V_1(2, 2))dt = \frac{\pi^M + \alpha\tilde{\pi}^D}{\alpha + r} > 0.$$

The firms share at $(2, 1)$ iff this maximizes their joint profits. Their joint profits under sharing are

$$V_J(2, 2) = V_1(2, 2) + V_2(2, 2) = 2\tilde{\pi}^D = \frac{2\pi^D}{r}$$

since when the firms share, the game reaches the history $(2, 2)$. Joint profits under no sharing are

$$V_1(2, 1, NS) + V_2(2, 1, NS) = \frac{\pi^M + 2\alpha\tilde{\pi}^D - c}{\alpha + r}.$$

We get $S \succ NS \iff$

$$\begin{aligned} 2\tilde{\pi}^D(\alpha + r) &> \pi^M + 2\alpha\tilde{\pi}^D - c \text{ or} \\ 2\pi^D + c &> \pi^M. \end{aligned}$$

This condition holds in the region, and **the firms share step 2 at** $(2, 1)$.

At the history $(1, 1)$, each firm has one success. There is no sharing decision to be made. The firms must, however, decide whether to invest to develop the second step. Assuming firm 1 invests, firm 2 will also invest if

$$V_2(1, 1) = \frac{\alpha V_2(2, 1) + \alpha V_2(1, 2) - c}{2\alpha + r} = \frac{\alpha V_J(2, 1) - c}{2\alpha + r} > 0$$

Since the firms share at $(2, 1)$, $V_J(2, 1) = 2\tilde{\pi}^D$. Substituting, we get

$$V_2(1, 1) = \frac{2\alpha\tilde{\pi}^D - c}{2\alpha + r} > 0 \quad (4)$$

This simplifies to

$$\pi^D > \frac{cr}{2\alpha}.$$

This condition holds in the region, so firm 2 invests. Hence, each firm invests at $(1, 1)$ if the other does. If firm 1 does not invest at $(1, 1)$, the new history is $(X, 1)$. Firm 2 invests if

$$V_2(X, 1) = \frac{\alpha V_2(X, 2) - c}{\alpha + r} = \frac{\alpha\tilde{\pi}^M - c}{\alpha + r} > 0 \quad (5)$$

where $V_2(X, 2) = \tilde{\pi}^M$ because at $(X, 2)$, **firm 2 produces output** as a monopolist. The condition simplifies to $\pi^M > \frac{cr}{\alpha}$. The condition holds because $\pi^M > \pi^D$ and in this region $\pi^D > \frac{cr}{\alpha}$. Hence, **firm 2 invests at $(X, 1)$** . It follows that **both firms invest at $(1, 1)$** .

At the history $(2, 0, NS)$, firm 1 produces output. The firms have decided not to share. Firm 2 invests iff

$$V_2(2, 0, NS) = \frac{\alpha V_2(2, 1) - c}{\alpha + r} > 0.$$

Since the lagging firm has no bargaining power, its earnings under sharing are the same as its earnings under no sharing at the history $(2, 1)$. The earnings under no sharing, $V_2(2, 1, NS)$, are given in (3). Substituting and rearranging gives us

$$\begin{aligned} V_2(2, 0, NS) &= \frac{\alpha^2\tilde{\pi}^D - c(2\alpha + r)}{(\alpha + r)^2} > 0 \text{ or} \\ \pi^D &> \frac{cr}{\alpha}\left(2 + \frac{r}{\alpha}\right). \end{aligned}$$

This condition fails in the region, so **firm 2 drops out at $(2, 0, NS)$** . (This result also follows from Lemma 1.)

To see whether the firms share step 1 at $(2, 0)$, we compare the joint profits under sharing with joint profits under no sharing. Joint profits under sharing are $V_J(2, 1) = 2\tilde{\pi}^D$ since if the firms share, the game reaches the history $(2, 1)$ and the firms share step 2. Joint profits under no sharing are $V_J(2, 0, NS) = V_1(2, X) = \tilde{\pi}^M$ since firm 2 drops out of the game if the firms do not share. In this region, we have that $\pi^M > 2\pi^D$. Hence, **the firms do not share at $(2, 0)$** . The lagging firm then drops out of the game.

Working backwards from either (2, 0) or (1, 1), we next consider the history (1, 0, NS). At this history, firm 1 has one success and firm 2 has no successes and the firms have decided not to share. Each firm must decide whether to invest. If firm 1 invests, then firm 2 also invests if

$$V_2(1, 0, NS) = \frac{\alpha V_2(1, 1) + \alpha V_2(2, 0) - c}{2\alpha + r} > 0 \quad (6)$$

We can substitute for $V_2(1, 1)$ from (4). Moreover, $V_2(2, 0) = 0$ since the firms do not share at (2, 0) and the lagging firm drops out. Substituting and simplifying, (6) becomes

$$\pi^D > \frac{cr}{\alpha} \left(\frac{3}{2} + \frac{r}{2\alpha} \right).$$

This holds in the region, so the lagging firm 2 invests at (1, 0, NS) if firm 1 does. It is straightforward to show that the leading firm 1 invests at (1, 0, NS) if firm 2 invests. If firm 2 does not invest, then the history becomes (1, X) and the leading firm invests as showed above. It follows that the leading firm invests at (1, 0, NS) whether or not the lagging firm invests. Thus, **both firms invest at (1, 0, NS)**.

To see whether the firms share step 1 at (1, 0), we compare joint profits under sharing with joint profits under no sharing. If the firms share, the game reaches the history (1, 1). Hence, joint profits are $V_J(1, 1)$. Joint profits under no sharing are

$$V_J(1, 0, NS) = \frac{\alpha V_J(2, 0) + \alpha V_J(1, 1) - 2c}{2\alpha + r} = \frac{\alpha \tilde{\pi}^M + \alpha V_J(1, 1) - 2c}{2\alpha + r}. \quad (7)$$

We have $NS \succ S \iff$

$$\alpha \tilde{\pi}^M + \alpha V_J(1, 1) - 2c > (2\alpha + r)V_J(1, 1)$$

Substituting for $V_J(1, 1) = 2V_2(1, 1)$ from (4) and simplifying, we have

$$\pi^M > 2\pi^D \left(\frac{2\alpha + 2r}{2\alpha + r} \right) + c \left(\frac{2r}{2\alpha + r} \right)$$

This inequality holds in the region, so **the firms do not share at (1, 0)**.

At the history (0, 0), assuming firm 2 invests, firm 1 will also invest if

$$V_1(0, 0) = \frac{\alpha V_1(1, 0, NS) + \alpha V_1(0, 1, NS) - c}{2\alpha + r} = \frac{\alpha V_J(1, 0, NS) - c}{2\alpha + r} > 0$$

Substituting using (7) and (4) and simplifying, this is

$$4\alpha\pi^D + (2\alpha + r)\pi^M > \frac{cr}{\alpha^2}(4\alpha + r)(2\alpha + r) + 2cr$$

Since $\pi^M > 2\pi^D$ in this region, the condition holds if

$$(8\alpha + 2r)\pi^D > \frac{cr}{\alpha^2}(4\alpha + r)(2\alpha + r) + 2cr$$

Since $\pi^D > \frac{cr}{\alpha}(\frac{3}{2} + \frac{r}{2\alpha})$ in this region, the condition holds if

$$(8\alpha + 2r)\frac{cr}{\alpha}(\frac{3}{2} + \frac{r}{2\alpha}) > \frac{cr}{\alpha^2}(4\alpha + r)(2\alpha + r) + 2cr.$$

This simplifies to

$$2\alpha(2\alpha + r) > 0$$

which always holds. Hence firm 1 invests at $(0, 0)$ if firm 2 invests.

Assuming firm 2 does not invest, the history becomes $(0, X)$. Firm 1 invests if

$$V_2(0, X) = \frac{\alpha V_2(1, X) - c}{\alpha + r} = \frac{\alpha(\frac{\alpha\tilde{\pi}^M - c}{\alpha + r}) - c}{\alpha + r} > 0$$

where we substituted for $V_2(1, X)$ using (5). Simplifying, we get

$$\pi^M > \frac{cr}{\alpha}(2 + \frac{r}{\alpha}). \quad (8)$$

In this region, we have that

$$\pi^M > 2\pi^D \text{ and } \pi^D > \frac{cr}{\alpha}(\frac{3}{2} + \frac{r}{2\alpha}).$$

These two conditions together imply that (8) holds. Hence, **firm 1 invests at $(0, X)$** . It follows that **both firms invest at $(0, 0)$** .

This completes the derivation of the equilibrium. The equilibrium is unique. The equilibrium is non-monotonic because the firms share at $(2, 1)$ but not at $(1, 0)$. The histories $(1, 0)$ and $(2, 1)$ are both reached on the equilibrium path, so the non-monotonicity arises on the equilibrium path.