

Massachusetts Institute of Technology
Department of Economics
Working Paper Series

Experimentation, Patents, and Innovation

**Daron Acemoglu
Kostas Bimpikis
Asuman Ozdaglar**

Working Paper 08-19
October 8, 2008

Room E52-251
50 Memorial Drive
Cambridge, MA 02142

This paper can be downloaded without charge from the
Social Science Research Network Paper Collection at
<http://ssrn.com/abstract=12817038>

Experimentation, Patents, and Innovation^{*}

Daron Acemoglu[†] Kostas Bimpikis[‡] Asuman Ozdaglar[§]

October 2008.

Abstract

This paper studies a simple model of experimentation and innovation. Our analysis suggests that patents may improve the allocation of resources by encouraging rapid experimentation and efficient *ex post* transfer of knowledge across firms. Each firm receives a private signal on the success probability of one of many potential research projects and decides when and which project to implement. A successful innovation can be copied by other firms. Symmetric equilibria (where actions do not depend on the identity of the firm) always involve delayed and staggered experimentation, whereas the optimal allocation never involves delays and may involve simultaneous rather than staggered experimentation. The social cost of insufficient experimentation can be arbitrarily large. Appropriately-designed patents can implement the socially optimal allocation (in all equilibria). In contrast to patents, subsidies to experimentation, research, or innovation cannot typically achieve this objective. We also show that when signal quality differs across firms, the equilibrium may involve a nonmonotonicity, whereby players with stronger signals may experiment after those with weaker signals. We show that in this more general environment patents again encourage experimentation and reduce delays.

Keywords: delay, experimentation, innovation, patents, research.

JEL Classification: O31, D83, D92.

^{*}We thank participants at the September 2009 workshop of the Toulouse Network on Information Technology for comments. Acemoglu also gratefully acknowledges financial support from the Toulouse Network on Information Technology.

[†]Dept. of Economics, Massachusetts Institute of Technology

[‡]Operations Research Center, Massachusetts Institute of Technology

[§]Dept. of Electrical Engineering and Computer Science, Massachusetts Institute of Technology

1 Introduction

Most modern societies provide intellectual property rights protection to innovators using a patent system. The main argument in favor of patents is that they encourage *ex ante* innovation by creating *ex post* monopoly rents (e.g., Arrow, 1962, Kitch, 1977, Reinganum, 1981, Tirole, 1988, Klemperer, 1990, Gilbert and Shapiro, 1990, Romer, 1990, Grossman and Helpman, 1991, Aghion and Howitt, 1992, Scotchmer, 1999, Gallini and Scotchmer, 2002). In this paper, we suggest an alternative (and complementary) social benefit to patents. We show that, under certain circumstances, patents encourage experimentation by potential innovators while still allowing socially beneficial transmission of knowledge across firms.

We construct a stylized model of experimentation and innovation. In our baseline game there is a large number of potential projects and N symmetric potential innovators (firms). Each firm receives a private signal on which of these projects is more likely to lead to an innovation and can decide to experiment with any of these projects at any point in time. A successful innovation is publicly observed and can be copied by any of the other potential innovators (for example, other firms can build on the knowledge revealed by the innovation in order to increase their own probability of success, but in the process capture some of the rents of this first innovator). The returns from the implementation of a successful innovation are nonincreasing in the number of firms implementing it. We provide an explicit characterization of the (subgame perfect or perfect Bayesian) equilibria of this dynamic game (both in discrete and continuous time). The symmetric equilibrium always features delayed and staggered experimentation. In particular, experimentation does not take place immediately and involves one firm experimenting before others (and the latter firms free-riding on the former's experimentation). In contrast, the optimal allocation never involves delays and may require simultaneous rather than staggered experimentation. The insufficient equilibrium incentives for experimentation may create a significant efficiency loss: the ratio of social surplus generated by the equilibrium relative to the optimal allocation can be arbitrarily small.

We next show that a simple patent system, where a copying firm has to make a prespecified payment to the innovator, can implement the optimal allocation. When the optimal allocation involves simultaneous experimentation, the patent system makes free-riding prohibitively costly and implements the optimal allocation as the unique equilibrium. When the optimal allocation involves staggered experimentation, the patent system plays a more subtle role. It permits *ex post* transmission of knowledge but still increases experimentation incentives to avoid delays. The patent system can achieve this because it generates "conditional" transfers. An innovator

receives a patent payment only when copied by other firms. Consequently, patents encourage one firm to experiment earlier than others, thus achieving rapid experimentation without sacrificing useful transfer of knowledge. Moreover, we show that patents can achieve this outcome in all equilibria. The fact that patents are particularly well designed to play this role is also highlighted by our result that while an appropriately-designed patent implements the optimal allocation in all equilibria, subsidies to experimentation, research, or innovation cannot achieve the same objective.

In our baseline model, both the optimal allocation and the symmetric equilibrium involve sequential experimentation. Inefficiency results from lack of sufficient experimentation or from delays. The structure of equilibria is richer when the strength (quality) of the signals received by potential innovators differs and is also private information. In this case, those with sufficiently strong signals will prefer not to copy successful innovations. With two firms or when the support of signals is such that no firm will have sufficiently strong signals, the “symmetric” equilibrium of this extended game (where strategies do not depend on the identity of the player) satisfies a *monotonicity* property, so that firms with stronger signals never act after firms with weaker signals. This ensures an efficient pattern of experimentation and efficient *ex post* transfer of knowledge. However, when there are more than two firms and sufficiently strong signals are possible, then the equilibrium may violate monotonicity. Interestingly, when this is the case patents are again potentially useful (though they cannot restore monotonicity without preventing (some) *ex post* transfer of knowledge).

In addition to the literature on patents mentioned above, a number of other works are related to our paper. First, ours is a simple model of (social) experimentation and shares a number of common features with recent work in this area (e.g., Bolton and Harris, 1999, 2000, and Keller, Rady and Cripps, 2005). These papers characterize equilibria of multi-agent two-armed bandit problems and show that there may be insufficient experimentation. The structure of the equilibrium is particularly simple in our model and can be characterized explicitly because all payoff-relevant uncertainty is revealed after a single successful experimentation. In addition, as discussed above, there is insufficient experimentation in our model as well, though this also takes a simple form: either there is free-riding by some firms reducing the amount of experimentation or experimentation is delayed. We also show that patent systems can increase experimentation incentives and implement the optimal allocation.

Second, the structure of equilibria with symmetric firms is reminiscent to equilibria in war of attrition games (e.g., Maynard Smith, 1974, Hendricks, Weiss and Wilson, 1988, Haigh and Cannings, 1989). War of attrition games have been used in the study of market exit by Fudenberg

and Tirole (1986) and Bulow and Klemperer (1999), research tournaments by Taylor (1995), and in auctions by Bulow and Klemperer (1994). In our symmetric model, as in symmetric wars of attrition, players choose the stochastic timing of their actions in such a way as to make other players indifferent and willing to mix over the timing of their own actions. The structure of equilibria and the optimal allocation is different, however, and the optimal allocation may involve either simultaneous experimentation by all players or staggered experimentation similar to those resulting in asymmetric equilibria. The novel beneficial role of patents in our model arises from their ability to implement such asymmetric equilibria.

Finally, the monotonicity property when the quality of signals differs across agents is similar to results in generalized wars of attrition (e.g., Fudenberg and Tirole, 1986, Bulow and Klemperer, 1994, 1999) and is also related to Gul and Lundholm's (1995) result on the clustering of actions in herding models. In the context of a standard herding model with endogenous timing, Gul and Lundholm construct an equilibrium in which agents with stronger signals act earlier than those with weaker signals, though the specifics of our model and analysis differs from these previous contributions. This is highlighted by the result that this monotonicity property does not hold in our model when there are more than two firms and the support of signals includes sufficiently strong signals so that some firms prefer not to copy successful experimentations.

The rest of the paper is organized as follows. In Section 2 we start with a discrete-time model with two symmetric firms. We characterize both asymmetric and symmetric equilibria in this model and provide explicit solutions when period length tends to zero. This motivates our continuous-time analysis, which is more tractable. We obtain the continuous-time model as the limit of the discrete-time model and provide explicit characterization of equilibria in Section 3. Section 4 extends these results to a setup with an arbitrary number of firms. Section 5 characterizes the optimal allocation and shows that the efficiency gap between the symmetric equilibrium and the optimal allocation can be arbitrarily large. The analysis in this section also demonstrates that an appropriately-designed patent system can implement the optimal allocation (in all equilibria). Section 6 extends the model to an environment in which signal quality differs across firms. It establishes the monotonicity property of equilibria when either there are only two firms or the support of signal distributions does not include sufficiently strong signals. Section 7 shows how this monotonicity property no longer holds when there are more than two firms and sufficiently strong signals are possible. Section 8 discusses the optimal allocation and the role of patents in this extended model. Section 9 concludes.

2 Two Symmetric Firms: Discrete Time

In this section, we start with a discrete-time model with two symmetric firms.

2.1 Environment

The economy consists of two research firms, each maximizing the present discounted value of profits. Let us denote the time interval between two consecutive periods by $\Delta > 0$. In what follows we will take Δ to be small.

Each firm can implement (“experiment with”) one of M potential innovation projects. Let us represent the set of projects by $\mathcal{M} = \{1, \dots, M\}$. Each firm receives a private (“positive”) signal $\varphi \in \mathcal{M}$ indicating the success potential of one of the projects. The unconditional probability that a project will be successful (when implemented) is small ($\simeq 0$). Conditional on the positive signal the success probability of a project is $p > 0$. We assume that the two firms always receive signals about different projects.¹ The success or failure of experimentation by a firm is publicly observed. When experimentation is successful, we refer to this as an “innovation”.

At each instant, a firm can choose one of three possible actions: (1) experiment with a project (in particular, with the project on which the firm has received a positive signal); (2) copy a successful project; (3) wait. Experimentation and copying are irreversible, so that a firm cannot then switch to implement a different project. In the context of research, this captures the fact that commitment of intellectual and financial resources to a specific research line or project is necessary for success. Copying of a successful project can be interpreted more broadly as using the information revealed by successful innovation or experimentation, so it does not need to correspond to the second firm replicating the exact same innovation (or product).²

Payoffs depend on the success of the project and whether the project is copied. During an interval of length Δ , the payoff to a firm that is the only one implementing a successful project is $\pi_1\Delta > 0$. In contrast, if a successful project is implemented by both firms, each receives $\pi_2\Delta > 0$.³ The payoff to an unsuccessful project is normalized to zero. Both firms discount the future at the common rate $r > 0$ (so that the discount factor per period is $e^{-r\Delta}$).

Until we introduce heterogeneity in success probabilities, we maintain the following assumption

¹This means that signals are not independent. If signals were independent, there would be a positive (but since M is large, small) probability that the two firms receive signals about the same project. Assuming that this event has zero probability simplifies notation.

²In line with this interpretation, we could also allow copying to be successful with some probability $\nu \in (p, 1]$. This added generality does not affect any of the main economic insights and we omit it.

³It will be evident from the analysis below that all of our results can be straightforwardly generalized to the case where an innovator receives payoff $\pi_2^{\text{first}}\Delta$ when copied, whereas the copier receives $\pi_2^{\text{second}}\Delta$. Since this has no major effect on the main economic insights and just adds notation, we do not pursue this generalization.

tion.⁴

Assumption 1

$$\pi_1 > \pi_2 > p\pi_1.$$

Let us also define the present discounted value of profits as

$$\Pi_j \equiv \frac{\pi_j}{r} \text{ for } j = 1, 2,$$

and for future reference, define

$$\beta \equiv \frac{\Pi_2}{\Pi_1}. \tag{1}$$

Clearly, $\beta \in (p, 1)$ in view of Assumption 1.

Now we are in a position to define strategies in this game. Let a history up to time t (where $t = k\Delta$ for some integer k) be denoted by h^t . The set of histories is denoted by \mathcal{H}^t . A strategy for a firm is a mapping from its signal, $\varphi \in \mathcal{M}$, and the history of the game up to time t , h^t , to the probability of experimentation at a given time interval and the distribution over projects. Thus the time t strategy can be written as

$$\sigma^t : \mathcal{M} \times \mathcal{H}^t \rightarrow [0, 1] \times \Delta(\mathcal{M}),$$

where $[0, 1]$ denotes the probability of implementing a project (either experimenting or copying) at time t and $\Delta(\mathcal{M})$ denotes the set of probability distributions over the set of projects, corresponding to the choice of project when the firm implements a project. The latter piece of generality is largely unnecessary (and will be omitted), since there will never be mixing over projects (a firm will either copy a successful project or experiment with the project for which it has received a positive signal). Here $\sigma^t(\varphi, h^t) = (0, \cdot)$ corresponds to waiting at time t and $\sigma^t(\varphi, h^t) = (1, j)$ corresponds to implementing project j at time t , which could be experimentation or copying of a successful project. Let us also denote the strategy of firm $i = 1, 2$ by $\sigma_i = \{\sigma_i^t\}_{t=0}^\infty$.

History up to time t can be summarized by two events $a^t \in \{0, 1\}$ denoting whether the other firm has experimented up to time t , $m^t \in \mathcal{M}$ denoting which project it has chosen in that case, and $s^t \in \{0, 1\}$ denoting whether this choice was successful. With a slight abuse of notation we will use both $\sigma^t(\varphi, h^t)$ and $\sigma^t(\varphi, a^t, m^t, s^t)$ to denote time t strategies. Although this is a game of incomplete information, the only source of asymmetric information is the identity of the project about which a firm has received a positive signal and no action other than the

⁴The structure of equilibria without this assumption is trivial as our analysis in Section 5 shows.

implementation of a project reveals information about this. In light of this, it is sufficient to focus on subgame perfect equilibria (rather than perfect Bayesian equilibria) in this game. A subgame perfect equilibrium (or simply *equilibrium*) is a strategy profile $(\hat{\sigma}_1, \hat{\sigma}_2)$ such that $\hat{\sigma}_i$ is the best response to $\hat{\sigma}_{-i}$ in all histories $h^t \in \mathcal{H}^t$ for $i = 1, 2$.

2.2 Asymmetric Equilibria

Even though firms are symmetric (in terms of their payoffs and information), there can be symmetric and asymmetric equilibria. Our main interest is with symmetric equilibria, where strategies are independent of the identity of the player. Nevertheless, it is convenient to start with asymmetric equilibria. These equilibria are somewhat less natural, because, as we will see, they involve one of the players never experimenting until the other one does.

In an asymmetric equilibrium, one of the firms, say 1, immediately attempts the project for which it received a signal. Firm 2 copies firm 1 in the next time period if the latter is successful and tries its own project otherwise. In terms of the notation above, this asymmetric equilibrium would involve

$$\hat{\sigma}_1^t(\varphi_1 = j, \cdot) = (1, j),$$

for $t = 0, \Delta, 2\Delta, \dots$. In words, this means that firm 1 chooses to experiment immediately (if it has not experimented yet until t) and experiments with the project on which it has received its signal, j . Firm 2, on the other hand, uses the strategy

$$\hat{\sigma}_2^t(\varphi_2 = j, a^t, m^t, s^t) = \begin{cases} (1, j') & \text{if } a^t = 1, m^t = j' \text{ and } s^t = 1, \\ (1, j) & \text{if } a^t = 1, m^t = j' \text{ and } s^t = 0, \\ (0, \cdot) & \text{if } a^t = 0, \end{cases}$$

for $t = 0, \Delta, 2\Delta, \dots$. The crucial feature highlighted by these strategies is that firm 2 never experiments until firm 1 does.

Using the same analysis as in the proof of Proposition 2 below, it is straightforward to verify that $\hat{\sigma}_2$ is a best response to $\hat{\sigma}_1$ provided that $\Delta < \Delta^* \equiv r^{-1} \log(\beta + 1 - p)$. What about $\hat{\sigma}_1$? Given $\hat{\sigma}_2$, suppose that the game has reached time t (where $t = k\Delta$ for $k \in \mathbb{N}$). If firm 1 now follows $\hat{\sigma}_1$, it will receive expected payoff

$$V[t] = p \left(e^{-rt} \pi_1 \Delta + e^{-r(t+\Delta)} \Pi_2 \right),$$

at time t , since its experimentation will be successful with probability p , yielding a profit of $\pi_1 \Delta$ during the first period following the success (equivalent to $e^{-rt} \pi_1 \Delta$ when discounted to time $t = 0$). Then according to $\hat{\sigma}_2$, firm 2 will copy the successful project and firm 1 will receive the present discounted value $e^{-r(t+\Delta)} \Pi_2$ from then on. If, at this point, firm 1 chooses not

to experiment, then the game proceeds to time $t + \Delta$, and according to the strategy profile $(\hat{\sigma}_1^{t+\Delta}, \hat{\sigma}_2^{t+\Delta})$, it will receive payoff equal to

$$V[t + \Delta] = e^{-r\Delta}V[t] < V[t].$$

Therefore this deviation is not profitable. This discussion establishes the following proposition (proof in the text). Throughout the paper, when there are two firms, we use the notation $\sim i$ to denote the firm $i' \neq i$.

Proposition 1 *Suppose that Assumption 1 holds and that $\Delta < \Delta^* \equiv r^{-1} \log(\beta + 1 - p)$. Then there exist two asymmetric equilibria. In each, one firm, $i = 1, 2$, tries its project with probability 1 immediately and the other, firm $\sim i$, never tries its project unless it observes the outcomes of the experimentation of firm i . Following experimentation by i , firm $\sim i$ copies it if successful and experiments with its own project otherwise.*

More formally, the two equilibria involve strategies of the form:

$$\sigma_i^t(\varphi_i = j, \cdot) = (1, j),$$

$$\sigma_{\sim i}^t(\varphi_{\sim i} = j, a^t, m_i^t, s^t) = \begin{cases} (1, j') & \text{if } a^t = 1, m_i^t = j' \text{ and } s^t = 1, \\ (1, j) & \text{if } a^t = 1, m_i^t = j' \text{ and } s^t = 0, \\ (0, \cdot) & \text{if } a^t = 0, \end{cases}$$

for $i = 1, 2$.

2.3 Symmetric Equilibria

Asymmetric equilibria explicitly condition on the identity of the firm: one of the firms, with label i , is treated differently than the firm with label $\sim i$. This has important payoff consequences. In particular, it can be verified easily that firm $\sim i$ has strictly greater payoffs in the equilibrium of Proposition 1 than firm i . In addition, as already noted in the previous section, asymmetric equilibria rely on the understanding by both firms that one of them will not experiment until the other one does. In this light, symmetric equilibria, where strategies are not conditioned on firms' "labels," and firms obtain the same equilibrium payoffs are more natural. In this subsection, we study such symmetric equilibria. We focus on the case where the time interval Δ is strictly positive but small.

As defined above a firm's strategy is a mapping from its information set to the probability of implementing a project. We refer to a strategy as pure if the experimentation probability at a given time t is either 0 or 1. That is, a pure strategy takes the form

$$\sigma^t : \mathcal{M} \times \mathcal{H}^t \rightarrow \{0, 1\} \times \mathcal{M}.$$

Our first result shows that for small Δ , there are no pure-strategy symmetric equilibria.

Proposition 2 *Suppose that Assumption 1 holds and that $\Delta < \Delta^* \equiv r^{-1} \log(\beta + 1 - p)$ (where recall that $\beta \equiv \Pi_2/\Pi_1$). Then there exist no symmetric pure-strategy equilibria.*

Proof. Suppose, to obtain a contradiction, that such an equilibrium σ^* exists. This implies that there exists some time t_0 and history h^{t_0} (with $t_0 = k\Delta$ for $k \in \mathbb{N}$) such that $\sigma^{t_0}(\varphi = j, h^{t_0}) = (1, j)$. Then following this history the payoff to both firms is

$$V[t_0 | \sigma^*, \sigma^*] = e^{-rt_0} p \Pi_1.$$

Now consider a deviation by firm 1 to σ' , which involves, after history h^{t_0} , waiting until date $t_0 + \Delta$, copying firm 2's project if successful, and experimenting with its own project otherwise. The payoff to this strategy is

$$V[t_0 | \sigma', \sigma^*] = e^{-r(t_0 + \Delta)} (p \Pi_2 + (1 - p) p \Pi_1),$$

since firm 2 experiments with probability 1 at time t_0 and is successful with probability p . Clearly, $V[t_0 | \sigma', \sigma^*] > V[t_0 | \sigma^*, \sigma^*]$ and there is a profitable deviation if

$$e^{-r(t_0 + \Delta)} (p \Pi_2 + (1 - p) p \Pi_1) > e^{-rt_0} p \Pi_1,$$

or if

$$\Delta < \Delta^* \equiv r^{-1} \log(\beta + 1 - p).$$

Here $\Delta^* > 0$ since, from Assumption 1, $\beta \equiv \Pi_2/\Pi_1 > p$. This establishes the existence of a profitable deviation and proves the proposition. ■

Proposition 2 is intuitive. Asymmetric equilibria involve one of the firms always waiting for the other one to experiment and receiving higher payoff. Intuitively, Proposition 2 implies that in symmetric equilibria both firms would like to be in the position of the firm receiving higher payoffs and thus delaying their own experimentation in order to benefit from that of the other firm. These incentives imply that no (symmetric) equilibrium can have immediate experimentation with probability 1 by either firm.

Proposition 2 also implies that all symmetric equilibria must involve mixed strategies. Moreover, any candidate equilibrium strategy must involve copying of a successful project in view of Assumption 1 and immediate experimentation when the other firm has experimented. Therefore, we can restrict attention to time t strategies of the form

$$\hat{\sigma}^t(\varphi = j, a^t, m^t, s^t) = \begin{cases} (1, j') & \text{if } a^t = 1, m^t = j' \text{ and } s^t = 1, \\ (1, j) & \text{if } a^t = 1, m^t = j' \text{ and } s^t = 0, \\ (q(t)\Delta, j) & \text{if } a^t = 0, \end{cases} \quad (2)$$

for $t = 0, \Delta, 2\Delta, \dots$, where $q(t)\Delta$ is the probability of experimenting at time t conditional on no experimentation by either firm up to time t (all such histories are identical, hence we write $q(t)$ instead of $q(h^t)$). Clearly, feasibility requires that $q(t)\Delta \leq 1$.

Next we derive an explicit characterization of the unique symmetric equilibrium as $\Delta \rightarrow 0$. In the text, we assume that firms use a constant probability of experimentation over time, i.e., $q(t) = q$ for all t (Proposition 3 relaxes this assumption and establishes uniqueness more generally). We consider a symmetric mixed-strategy equilibrium σ^* and suppose that the game has reached time t without experimentation. Let $v_w[t | \Delta]$ and $v_e[t | \Delta]$ denote the time t continuation payoffs to firm i when firm $\sim i$ plays σ^* and firm i chooses to wait or to experiment (and period length is Δ).⁵ For a mixed-strategy equilibrium to exist, we need

$$v_w[t | \Delta] = v_e[t | \Delta]. \quad (3)$$

The proof of Proposition 3 below shows that all symmetric equilibria involve mixing after any history h^t (with no experimentation up to t), i.e., equation (3) holds for all such h^t . Therefore, it suffices to characterize σ^* such that (3) holds. First, consider firm i 's payoffs from experimenting:

$$v_e[t | \Delta] = q\Delta p\Pi_1 + (1 - q\Delta)p\Pi_2, \quad (4)$$

since in this case firm i is successful with probability p and receives continuation value Π_1 if firm $\sim i$ has also experimented during the same time interval (probability $q\Delta$), and it receives Π_2 otherwise (probability $1 - q\Delta$).

Similarly, its payoff from waiting is

$$v_w[t | \Delta] = e^{-r\Delta} \left(q\Delta (p\Pi_2 + (1 - p)p\Pi_1) + (1 - q\Delta)v_w[t + \Delta | \Delta] \right), \quad (5)$$

where firm i receives no payoff today and with probability $q\Delta$, firm $\sim i$ experiments, in which case firm i copies if the experimentation is successful and experiments with its own project otherwise, with expected continuation return $p\Pi_2 + (1 - p)p\Pi_1$. With probability $1 - q\Delta$, firm $\sim i$ does not experiment, and firm i then receives $v_w[t + \Delta | \Delta]$. Adding and subtracting $v_w[t + \Delta | \Delta]$ from the left-hand side of (5) and rearranging, we obtain

$$v_w[t + \Delta | \Delta] (1 - (1 - q\Delta) e^{-r\Delta}) - (v_w[t + \Delta | \Delta] - v_w[t | \Delta]) = e^{-r\Delta} q\Delta (p\Pi_2 + (1 - p)p\Pi_1).$$

Dividing both sides by Δ and taking the limit as $\Delta \rightarrow 0$ yields

$$\lim_{\Delta \rightarrow 0} v_w[t + \Delta | \Delta] - \frac{1}{r + q} \lim_{\Delta \rightarrow 0} \left(\frac{v_w[t + \Delta | \Delta] - v_w[t | \Delta]}{\Delta} \right) = \frac{q}{r + q} [p\Pi_2 + (1 - p)p\Pi_1]. \quad (6)$$

⁵Here we use v , since V denotes the value discounted back to $t = 0$.

From equation (4), we see that $v_e[t | \Delta]$ does not depend on t . Since equation (3) holds for all h^t (thus for all t), we have $v_w[t | \Delta] = v_e[t | \Delta]$ and $v_w[t + \Delta | \Delta] = v_e[t + \Delta | \Delta]$, implying that $v_w[t + \Delta | \Delta] = v_w[t | \Delta]$. Therefore, the second term on the left-hand side of (6) must be equal to zero. Moreover, taking the limit as $\Delta \rightarrow 0$ in (4), we obtain

$$\lim_{\Delta \rightarrow 0} v_w[t + \Delta | \Delta] = \lim_{\Delta \rightarrow 0} v_e[t | \Delta] = p\Pi_2.$$

Combined with (6), this yields

$$q(t) = q^* \equiv \frac{r\beta}{1-p} \text{ for all } t, \quad (7)$$

where recall, from (1), that $\beta \equiv \Pi_2/\Pi_1$.

The next proposition relaxes the assumption that $q(t)$ is constant for all t and shows that this is indeed the unique symmetric equilibrium.

Proposition 3 *Suppose that Assumption 1 holds and $\Delta \rightarrow 0$. Then there exists a unique symmetric equilibrium. In this equilibrium, both firms use the mixed strategy $\hat{\sigma}$ as given in (2) with $q(t) = q^*$ as in (7).*

Proof. We first show that any symmetric equilibrium must involve mixing after any history $h^t \in \mathcal{H}^t$ along which there has been no experimentation. The argument in the proof of Proposition 2 establishes that after any such history h^t , there cannot be experimentation with probability 1. We next show that there is positive probability of experimentation at time $t = 0$. First note that the equilibrium-path value to a firm, V^* , (discounted back to time $t = 0$), satisfies $V^* \geq p\Pi_2$, since each firm can guarantee this by experimenting at time $t = 0$. This implies that in any equilibrium there must exist some time T such that after time T there is positive probability of experimentation and innovation. Now to obtain a contradiction, suppose that $T > 0$. By the argument preceding the proposition, $\lim_{\Delta \rightarrow 0} V_e[T | \Delta] = e^{-rT}p\Pi_2$, and therefore, in any mixed-strategy equilibrium, $V^*[T | \Delta] \rightarrow e^{-rT}p\Pi_2$. However, for $T > 0$ this is strictly less than $V^* \geq p\Pi_2$, yielding a contradiction and establishing the desired result. The same argument also establishes that there cannot exist any time interval (T, T') , with $T' > T$, along which there is no mixing.

Hence, along any history h^t where there has not been an experimentation, both firms must be indifferent between waiting and experimenting. This implies that (3) must hold for all t . Let $q(t)\Delta$ denote the probability of experimentation at time t . Firm i 's payoff for experimenting at time t is given by an expression similar to equation (4),

$$v_e[t | \Delta] = q(t)\Delta p\Pi_1 + (1 - q(t)\Delta)p\Pi_2. \quad (8)$$

We next show that the probability of experimentation $q(t)$ in a symmetric equilibrium is a continuous function of t . Suppose that $q(t)$ is not continuous at some $\bar{t} \geq 0$. If $q(\bar{t}) < q(\bar{t}+)$ (where $q(\bar{t}+) \equiv \lim_{t \downarrow \bar{t}} q(t)$), it then follows from (8) and Assumption 1 that $v_e[\bar{t} | \Delta] < v_e[\bar{t}+ | \Delta]$. This implies that firm i has an incentive to delay experimentation at time \bar{t} . But this contradicts the fact that the symmetric equilibrium must involve mixing at all such t . Similarly, if $q(\bar{t}) > q(\bar{t}+)$, we have $v_e[\bar{t} | \Delta] > v_e[\bar{t}+ | \Delta]$, implying that firm i will experiment with probability 1 at time \bar{t} , again yielding a contradiction. This establishes that $q(t)$ is a continuous function of t .

A derivation similar to that preceding the proposition then shows that equation (6) holds when q is replaced by $q(t)$. In particular,

$$\lim_{\Delta \rightarrow 0} v_w[t + \Delta | \Delta] - \frac{1}{r + q(t)} \lim_{\Delta \rightarrow 0} \left(\frac{v_w[t + \Delta | \Delta] - v_w[t | \Delta]}{\Delta} \right) = \frac{q(t)}{r + q(t)} [p\Pi_2 + (1 - p)p\Pi_1]. \quad (9)$$

Since $v_w[t | \Delta] = v_e[t | \Delta]$ and $v_w[t + \Delta | \Delta] = v_e[t + \Delta | \Delta]$, we can write the second term on the left-hand side of equation (9) as

$$\begin{aligned} \lim_{\Delta \rightarrow 0} \left(\frac{v_w[t + \Delta | \Delta] - v_w[t | \Delta]}{\Delta} \right) &= \lim_{\Delta \rightarrow 0} \left(\frac{v_e[t + \Delta | \Delta] - v_e[t | \Delta]}{\Delta} \right) \\ &= \lim_{\Delta \rightarrow 0} \left(q(t + \Delta) - q(t) \right) \Delta p(\Pi_1 - \Pi_2) \\ &= 0, \end{aligned}$$

where the second equality follows from equation (8) and the third equality holds by the continuity of the experimentation probability $q(t)$. Substituting for

$$\lim_{\Delta \rightarrow 0} v_w[t + \Delta | \Delta] = \lim_{\Delta \rightarrow 0} v_e[t | \Delta] = p\Pi_2,$$

in equation (9) and solving for $q(t)$ yields $q(t) = q^*$ as in (7), completing the proof. ■

Proposition 3 characterizes the unique mixed strategy equilibrium when period length Δ is small. The analysis is in discrete time. The limit where $\Delta \rightarrow 0$, which Proposition 3 shows is well behaved, should also correspond to the (symmetric) equilibrium of the same model set up directly in continuous time. We establish this in the next section and subsequently work directly with the continuous-time model, which is both more economical and slightly more tractable.

3 Two Symmetric Firms: Continuous Time

The model is identical to that introduced in the previous section, except that time is continuous and we work directly with flow rates of experimentation. Strategies and equilibria are defined

similarly and Assumption 1 still applies. This implies that in any equilibrium a firm that has not experimented yet will copy a successful innovation. Using this observation, we can represent strategies by a function $\lambda : \mathbb{R}_+ \rightarrow \bar{\mathbb{R}}_+$ (where $\bar{\mathbb{R}}_+ \equiv \mathbb{R}_+ \cup \{+\infty\}$) specifying the flow rate of experimentation at each date until there is experimentation by one of the players. Waiting with probability 1 at time t corresponds to $\lambda(t) = 0$, while experimenting with probability 1 at some time t corresponds to $\lambda(t) = +\infty$. The flow rate of experimentation λ induces a stochastic distribution of “stopping time,” which we denote by τ . The stopping time τ designates the probability distribution that experimentation will happen at any time $t \in \bar{\mathbb{R}}_+$ conditional on the other player not having experimented until then. A pure strategy simply specifies $\tau \in \bar{\mathbb{R}}_+$. For example, the strategy of experimenting immediately is $\tau = 0$, whereas that of waiting for the other firm’s experimentation is represented by $\tau = +\infty$. The τ notation is convenient to use for the next two propositions, while in characterizing the structure of equilibria we need to use λ (thus justifying the introduction of both notations).⁶

Proposition 4 *Suppose Assumption 1 holds. Then there exist two asymmetric equilibria. In each equilibrium, $\tau_i = 0$ and $\tau_{\sim i} = +\infty$ for $i = 1, 2$.*

Proof. The proof follows from the same argument as in subsection 2.2 and is omitted. ■

Our next result is the parallel of Proposition 2 and shows that there are no symmetric pure-strategy equilibria in the continuous-time model.

Proposition 5 *Suppose Assumption 1 holds. Then there exist no symmetric pure-strategy equilibria.*

Proof. Suppose, to obtain a contradiction, that a symmetric pure-strategy equilibrium exists. Then $\tau^* = t \in \mathbb{R}_+$ for $i = 1, 2$, yielding payoff

$$V(\tau^*, \tau^*) = e^{-rt} p \Pi_1$$

to both players. Now consider a deviation $\tau' > t$ for one of the firms, which involves waiting for a time interval ϵ and copying a successful innovation if there is such an innovation during this time interval. As $\epsilon \rightarrow 0$, this strategy gives the deviating firm payoff equal to

$$V(\tau', \tau^*) = \lim_{\epsilon \downarrow 0} e^{-r(t+\epsilon)} [p \Pi_2 + (1-p)p \Pi_1].$$

⁶Here we can be more formal and follow Simon and Stinchcombe’s (1989) formal model of continuous-time games with jumps. This amounts to defining an extended strategy space, where stopping times are defined for all $t \in \bar{\mathbb{R}}_+$ and also for $t+$ for any $t \in \mathbb{R}_+$. In other words, strategies will be piecewise continuous and right continuous functions of time, so that a jump immediately following some time $t \in \mathbb{R}_+$ is well defined. Throughout, we do allow such jumps, but do not introduce the additional notation, since this is not necessary for any of the main economic insights or proofs.

Assumption 1 implies that $V(\tau', \tau^*) > V(\tau^*, \tau^*)$, establishing the result. ■

As in the discrete-time model, we will next show that there exists a (unique) symmetric mixed-strategy equilibrium. In what follows, instead of working with the stopping time τ , it is more convenient to work directly with $\lambda(t)$, which designates the flow rate of experimentation at time t . The next lemma shows that symmetric equilibria must involve mixing on all $t \in \mathbb{R}_+$ and will be used in the characterization of mixed-strategy equilibria.

Lemma 1 *The support of mixed strategy equilibria is \mathbb{R}_+ .*

Proof. The proof comprises three steps.

First, we show that $t = 0$ belongs to the support of mixing time (so that there is no time interval with zero probability of experimentation). Suppose, to obtain a contradiction, that $t_1 = \inf\{t: \lambda(t) > 0\} > 0$. Then with the same argument as in the proof of Proposition 3 (in particular because experimenting after t_1 is in the support of the mixed-strategy equilibrium), equilibrium payoffs must be

$$V_1 = e^{-rt_1} p\Pi_2.$$

Now consider deviation where firm i chooses $\lambda(0) = +\infty$. This has payoff

$$V_0 = p\Pi_2 > V_1$$

for any $t_1 > 0$, yielding a contradiction.

Second, we show that there does not exist $T < \infty$ such that the support of the stopping time τ (induced by λ) is within $[0, T]$. Suppose not, then it implies that there exists $t \in [0, T]$ such that $\lambda(t) = +\infty$ and let $t_1 = \inf\{t: \lambda(t) = +\infty\}$. This implies that the payoff to both firms once the game reaches time t_1 without experimentation (which has positive probability since $t_1 = \inf\{t: \lambda(t) = +\infty\}$) is

$$V(\tau = t_1) = e^{-rt_1} p\Pi_2$$

(where $V(\tau = t)$, or $V(t)$, denotes present discounted value as a function of experimentation time; this is different than $V[t]$, which referred to the value at time t in subsection 2.2). Now consider a deviation by firm i to strategy τ' , which involves waiting for $\epsilon > 0$ after the game has reached t_1 and copying a successful project by firm $\sim i$ (if there is such a success). This has payoff

$$V(\tau') = e^{-r(t+\epsilon)} [p\Pi_2 + (1-p)p\Pi_1]$$

since firm $\sim i$ is still $\lambda(t_1) = +\infty$ and will thus experiment with probability 1 at t_1 . Assumption 1 implies that $V(\tau')$ is strictly greater than $V(\tau = t_1)$ for ϵ sufficiently small.

Finally, we show that $\lambda(t) > 0$ for all t . Again suppose, to obtain a contradiction, that there exist t_1 and $t_2 > t_1$ such that $\lambda(t) = 0$ for $t \in (t_1, t_2)$. Then, with the same argument as in the first part, the payoff from the candidate equilibrium strategy τ to firm i conditional on no experimentation until t_1 is

$$V(\tau) = e^{-rt_2} p \Pi_2.$$

However, deviating and choosing $\tau' = t_1$ yields

$$V(\tau' = t_1) = e^{-rt_1} p \Pi_2 > V(\tau).$$

This contradiction completes the proof of the lemma. ■

Lemma 1 implies that in all symmetric equilibria there will be mixing at all times (until there is experimentation). Using this observation, Proposition 6 characterizes a unique symmetric equilibrium. Let us illustrate the reasoning here by assuming that firms use a constant flow rate of experimentation (the proof of Proposition 6 relaxes this assumption). In particular, suppose that firm $\sim i$ innovates at the flow rate λ for all $t \in \mathbb{R}_+$. Then the value of innovating at time t (i.e., choosing $\tau = t$) for firm i is

$$V(t) = \int_0^t \lambda e^{-\lambda z} e^{-rz} [p \Pi_2 + (1-p) p \Pi_1] dz + e^{-\lambda t} e^{-rt} p \Pi_2. \quad (10)$$

This expression uses the fact that when firm $\sim i$ is experimenting at the flow rate λ , the timing of its experimentation has an exponential distribution, with density $\lambda e^{-\lambda t}$. Then the first term in (10) is the expected discounted value from the experimentation of firm $\sim i$ between 0 and t (again taking into account that following an experimentation, a successful innovation will be copied, with continuation value $p \Pi_2 + (1-p) p \Pi_1$). The second term is the probability that firm $\sim i$ does not experiment until t , which, given the exponential distribution, is equal to $e^{-\lambda t}$, multiplied by the expected discounted value to firm i when it is the first to experiment at time t (given by $e^{-rt} p \Pi_2$).

Lemma 1 implies that $V(t)$ must be constant in t for all $t \in \mathbb{R}_+$. Therefore, its derivative $V'(t)$ must be equal to zero for all t in any symmetric equilibrium implying that

$$V'(t) = \lambda e^{-\lambda t} e^{-rt} [p \Pi_2 + (1-p) p \Pi_1] - (r + \lambda) e^{-\lambda t} e^{-rt} p \Pi_2 = 0, \quad (11)$$

for all t . This equation has a unique solution:

$$\lambda^* \equiv \frac{r\beta}{1-p} \text{ for all } t. \quad (12)$$

The next proposition shows that this result also holds when both firms can use time-varying experimentation rates.

Proposition 6 *Suppose Assumption 1 holds. Then there exists a unique symmetric equilibrium. This equilibrium involves both firms using a constant flow rate of experimentation λ^* as given by (12). Firm $\sim i$, firm i immediately copies a successful innovation by firm $\sim i$ and experiments if firm $\sim i$ experiments unsuccessfully.*

Proof. Suppose that firm $\sim i$ experiments at the flow rate $\lambda(t)$ at time $t \in \mathbb{R}_+$. Let us define

$$m(t) \equiv \int_0^t \lambda(z) dz. \quad (13)$$

Then the equivalent of (10) is

$$V(t) = \int_0^t \lambda(z) e^{-m(z)} e^{-rz} [p\Pi_2 + (1-p)p\Pi_1] dz + e^{-m(t)} e^{-rt} p\Pi_2. \quad (14)$$

Here $\int_{t_1}^{t_2} \lambda(z) e^{-m(z)} dz$ is the probability that firm $\sim i$ (using strategy λ) will experiment between times t_1 and t_2 , and $e^{-m(t)} = 1 - \int_0^t \lambda(z) e^{-m(z)} dz$ is the probability that $\sim i$ has not experimented before time t . Thus the first term is the expected discounted value from the experimentation of firm $\sim i$ between 0 and t (discounted and multiplied by the probability of this event). The second term is again the probability that firm $\sim i$ does not experiment until t multiplied by the expected discounted value to firm i when it is the first to experiment at time t (given by $e^{-rt} p\Pi_2$).

Lemma 1 implies that $V(t)$ must be constant in t for all $t \in \mathbb{R}_+$. Since $V(t)$ is differentiable in t , this implies that its derivative $V'(t)$ must be equal to zero for all t . Therefore,

$$\begin{aligned} V'(t) &= \lambda(t) e^{-m(t)} e^{-rt} [p\Pi_2 + (1-p)p\Pi_1] - (r + m'(t)) e^{-m(t)} e^{-rt} p\Pi_2 \\ &= 0 \text{ for all } t. \end{aligned}$$

Moreover, note that $m(t)$ is differentiable and $m'(t) = \lambda(t)$. Therefore, this equation is equivalent to

$$\lambda(t) [p\Pi_2 + (1-p)p\Pi_1] = (r + \lambda(t)) p\Pi_2 \text{ for all } t. \quad (15)$$

The unique solution to (15) is (12), establishing the uniqueness of the symmetric equilibrium without restricting strategies to constant flow rates. ■

This proposition contains exactly the same economics as Proposition 3 derived in the previous section. The flow rate of innovation λ^* in (12) is also clearly identical to q^* in (7). This establishes formally that the limit of the discrete-time model and the continuous-time model give the same economic and mathematical answers. In what follows, we will use the continuous-time model, since it is more tractable and necessitates less notation.

4 Multiple Firms

Let us now suppose that there are N firms, each of which receives a positive signal about one of the projects. The probability that the project that has received a positive signal will succeed is still p and each firm receives a signal about a different project. Let π_n denote the flow payoff from a project that is implemented by n other firms and define

$$\Pi_n \equiv \frac{\pi_n}{r}.$$

Once again, $\beta \equiv \Pi_2/\Pi_1$ as specified in (1) and Assumption 1 holds, so that $\beta > p$.

The following proposition is established using similar arguments to those in the previous two sections and its proof is omitted.

Proposition 7 *Suppose that Assumption 1 holds and that there are $N \geq 2$ firms. Then there exist no symmetric pure-strategy equilibria. Moreover the support of the mixed-strategy equilibria is \mathbb{R}_+ .*

It is also straightforward to show that there exist asymmetric pure-strategy equilibria. For example, when $\Pi_N/\Pi_1 > p$, it is an equilibrium for firm 1 to experiment and the remaining $N - 1$ to copy if this firm is successful. If it is unsuccessful, then firm 2 experiments and so on.

As in the previous two sections, symmetric equilibria are of greater interest. To characterize the structure of symmetric equilibria, let us first suppose that

$$\Pi_n = \Pi_2 \text{ for all } n \geq 2 \tag{16}$$

and also to simplify the discussion, focus on symmetric equilibria with constant flow rates. In particular, let the rate of experimentation when there are $n \geq 2$ firms be λ_n . Consider a subgame starting at time t_0 with n firms that have not yet experimented (and all previous, $N - n$, experiments have been unsuccessful). Then the continuation value of firm i (from time t_0 onwards) when it chooses to experiment with probability 1 at time $t_0 + t$ is

$$v_n(t) = \int_{t_0}^{t_0+t} \lambda_n(n-1) e^{-\lambda_n(n-1)(z-t_0)} e^{-r(z-t_0)} [p\Pi_2 + (1-p)v_{n-1}] dz + e^{-\lambda_n(n-1)t} e^{-rt} p\Pi_2, \tag{17}$$

where v_{n-1} is the maximum value that the firm can obtain when there are $n - 1$ firms that have not yet experimented (where we again use v since this expression refers to the continuation value from time t_0 onwards). Intuitively, $\lambda_n(n-1) e^{-\lambda_n(n-1)(z-t_0)}$ is the density at which one of the $n - 1$ other firms mixing at the rate λ_n will experiment at time $z \in (t_0, t_0 + t)$. When this happens, it is successful with probability p and will be copied by all other firms, and each will

receive a value of $e^{-r(z-t_0)}\Pi_2$ (discounted back to t_0). If it is unsuccessful (probability $1-p$), the number of remaining firms is $n-1$, and this gives a value of v_{n-1} . If no firm experiments until time t , firm i chooses to experiment at this point and receives $e^{-rt}p\Pi_2$. The probability of this event is $1 - \int_{t_0}^{t_0+t} \lambda_n (n-1) e^{-\lambda_n(n-1)(z-t_0)} dz = e^{-\lambda_n(n-1)t}$. As usual, in a mixed strategy equilibrium, $v_n(t)$ needs to be independent of t and moreover, it is clearly differentiable in t . So its derivative must be equal to zero. This implies

$$\lambda_n (n-1) [p\Pi_2 + (1-p)v_{n-1}] = (\lambda_n (n-1) + r)p\Pi_2. \quad (18)$$

Proposition 7 implies that there has to be mixing in all histories, thus

$$v_n = p\Pi_2 \text{ for all } n \geq 2. \quad (19)$$

Intuitively, mixing implies that the firm is indifferent between experimentation and waiting, and thus its continuation payoff must be the same as the payoff from experimenting immediately, which is $p\Pi_2$. Combining (18) and (19) yields

$$\lambda_n = \frac{r\Pi_2}{(1-p)(n-1)\Pi_1}. \quad (20)$$

This derivation implies that in the economy with N firms, each firm starts mixing at the flow rate λ_N . Following an unsuccessful experimentation, they increase their flow rate of experimentation to λ_{N-1} , and so on.

The derivation leading up to (20) easily generalizes when we relax (16). To demonstrate this, let us relax (16) and instead strengthen Assumption 1 to:

Assumption 2

$$\Pi_n > p\Pi_1 \text{ for all } n.$$

The value of experimenting at time t (starting with n firms) is now given by a generalization of (17):

$$v_n(t) = \int_{t_0}^{t_0+t} \lambda_n (n-1) e^{-\lambda_n(n-1)(z-t_0)} e^{-r(z-t_0)} [p\Pi_n + (1-p)v_{n-1}] dz + e^{-\lambda_n(n-1)t} e^{-rt} p\Pi_n.$$

Again differentiating this expression with respect to t and setting the derivative equal to zero gives the equivalent indifference condition to (18) as

$$\lambda_n (n-1) [p\Pi_n + (1-p)v_{n-1}] = (\lambda_n (n-1) + r)p\Pi_n. \quad (21)$$

for $n = 2, \dots, N$. In addition, we still have

$$V_n = p\Pi_n \text{ for all } n \geq 2.$$

Combining this with (21), we obtain

$$\lambda_n = \frac{r\Pi_n}{(1-p)(n-1)\Pi_{n-1}} \text{ for } n = 2, \dots, N, \quad (22)$$

and let us adopt the convention that $\lambda_1 = +\infty$.

This derivation establishes the following proposition.

Proposition 8 *Suppose that Assumption 2 holds. Then there exists a unique symmetric equilibrium. In this equilibrium, when there are $n = 1, 2, \dots, N$ firms that have not yet experimented, each experiments at the constant flow rate λ_n as given by (22). A successful innovation is immediately copied by all remaining firms. An unsuccessful experimentation starting with $n \geq 3$ firms is followed by all remaining firms experimenting at the flow rate λ_{n-1} .*

An interesting feature of Proposition 8 is that after an unsuccessful experimentation, the probability of further experimentation may decline. Whether this is the case or not depends on how fast Π_n decreases in n .

5 Patents and Optimal Allocations

The analysis so far has established that both in discrete and continuous time, symmetric equilibria involve mixed strategies, potential delays, and also staggered experimentation (meaning that with probability 1, one of the firms will experiment before others). Asymmetric equilibria avoid delays, but also feature staggered experimentation. Moreover, they are less natural, because they involve one of the firms never acting (experimenting) until the other one does and also because they give potentially very different payoffs to different firms. In this section, we first establish the inefficiency of (symmetric) equilibria. We then suggest that an appropriately-designed patent system can implement optimal allocations. While all of the results in the section hold for $N \geq 2$ firms, we focus on the case with two firms to simplify notation.

5.1 Welfare

It is straightforward to see that symmetric equilibria are Pareto suboptimal. Suppose that there exists a social planner that can decide the experimentation time for each firm. Suppose also that the social planner would like to maximize the sum of the present discounted values of the two firms. Clearly, in practice an optimal allocation (and thus the objective function of the social planner) may also take into account the implications of these innovations on consumers. However, we have not so far specified how consumer welfare is affected by the replication of

successful innovations versus new innovations. Therefore, in what follows, we focus on optimal allocations from the viewpoint of firms. This would also be the optimal allocation taking consumer welfare into account when consumer surpluses from a new innovation and from a successful innovation implemented by two firms are proportional to Π_1 and $2\Pi_2$, respectively. If we take the differential consumer surpluses created by these innovations into account, this would only affect the thresholds provided below, and for completeness, we also indicate what these alternative thresholds would be (see, in particular, footnote 7).

The social planner could adopt one of two strategies:

1. *Staggered experimentation*: this would involve having one of the firms experiment at $t = 0$; if it is successful, then the other firm would copy the innovation, and otherwise the other firm would experiment immediately. Denote the surplus generated by this strategy by S_1^P .
2. *Simultaneous experimentation*: this would involve having both firms experiment immediately at $t = 0$. Denote the surplus generated by this strategy by S_2^P .

It is clear that no other strategy could be optimal for the planner. Moreover, both in discrete time as $\Delta \rightarrow 0$ and in continuous time, S_1^P and S_2^P have simple expressions. In particular,

$$S_1^P = 2p\Pi_2 + (1 - p)p\Pi_1. \quad (23)$$

Intuitively, one of the firms experiments first and is successful with probability p . When this happens, the other firm copies a successful innovation, with total payoff $2\Pi_2$. With the complementary probability, $1 - p$, the first firm is unsuccessful, and the second firm experiments independently, with expected payoff $p\Pi_1$. Both in continuous time and in discrete time as $\Delta \rightarrow 0$, these payoffs occur immediately after the first experimentation and thus are not discounted.

The alternative is to have both firms experiment immediately, which generates expected surplus

$$S_2^P = 2p\Pi_1. \quad (24)$$

The comparison of S_1^P and S_2^P implies that simultaneous experimentation by both firms is optimal when $2\beta < 1 + p$. In contrast, when $2\beta > 1 + p$, the optimal allocation involves one of the firms experimenting first, and the second firm copying successful innovations. This is stated in the next proposition (proof in the text).⁷

Proposition 9 *Suppose that*

$$2\beta \geq 1 + p, \quad (25)$$

⁷If consumer surpluses from a new innovation and from the two firms implementing the same project were,

then the optimal allocation involves staggered experimentation, that is, experimentation by one firm and copying of successful innovations. If (25) does not hold, then the optimal allocation involves immediate experimentation by both firms. When $2\beta = 1 + p$, both staggered experimentation and immediate experimentation are socially optimal.

Let us now compare this to the equilibria characterized so far. Clearly, asymmetric equilibria are identical to the first strategy of the planner and thus generate surplus S_1^P (recall subsection 2.2.). In contrast, the (unique) symmetric equilibrium generates social surplus

$$\begin{aligned} S^E &= \int_0^\infty 2\lambda^* e^{-(2\lambda^*+r)t} [2p\Pi_2 + (1-p)p\Pi_1] dt \\ &= \frac{2\lambda^*}{2\lambda^*+r} [2p\Pi_2 + (1-p)p\Pi_1], \end{aligned} \quad (26)$$

where λ^* is the (constant) equilibrium flow rate of experimentation given by (12). The first line of (26) applies because the time of first experimentation corresponds to the first realization of one of two random variables, both with an exponential distribution with parameter λ^* and time is discounted at the rate r . If the first experimentation is successful, which has probability p , surplus is equal to $2\Pi_2$, and otherwise (with probability $1-p$), the second firm also experiments, with expected payoff $p\Pi_1$. The second line is obtained by solving the integral and substituting for (23). It is also straightforward to verify that the second line of (26) applies in the discrete-time model with $\Delta \rightarrow 0$ (since (12) is identical to (7)).

A straightforward comparison shows that S^E is always (strictly) less than S_1^P . Therefore, the unique symmetric equilibrium is always inefficient. Moreover, this inefficiency can be quantified in a simple manner. Let $S^P = \max\{S_1^P, S_2^P\}$ and consider the ratio of equilibrium social surplus to the social surplus in the optimal allocation as a measure of inefficiency:

$$\mathfrak{s} \equiv \frac{S^E}{S^P}.$$

Naturally, the lower is \mathfrak{s} the more inefficient is the equilibrium.

Clearly, $\mathfrak{s} < 1$, so that the equilibrium is always inefficient as stated above. More specifically, let us first suppose that (25) holds. Then, the source of inefficiency is delayed experimentation.

respectively, C_1 and $2C_2$, then we would have

$$\begin{aligned} S_1^P &= 2p(\Pi_2 + C_2) + (1-p)p(\Pi_1 + C_1) \\ S_2^P &= 2p(\Pi_1 + C_1). \end{aligned}$$

Denoting $\gamma \equiv (\Pi_2 + C_2) / (\Pi_1 + C_1)$, it is then clear that condition (25) would be replaced by $2\gamma \geq 1 + p$ and the rest of the analysis would remain unchanged. As noted in the beginning of this section, if $C_2 = \kappa\Pi_2$ and $C_1 = \kappa\Pi_1$, then this condition would be identical to (25).

In this case,

$$\begin{aligned}\mathfrak{s} &= \frac{S^E}{S_1^P} \\ &= \frac{2\lambda^*}{2\lambda^* + r} = \frac{2\beta}{2\beta + 1 - p},\end{aligned}$$

where the last equality simply uses (12). It is clear that \mathfrak{s} is minimized, for given p , as $\beta = (1 + p)/2$ (its lower bound given (25)). In that case, we have

$$\mathfrak{s} = \frac{1 + p}{2}.$$

In addition, as $p \downarrow 0$, \mathfrak{s} can be as low as $1/2$.

Next consider the case where (25) does not hold. Then

$$\begin{aligned}\mathfrak{s} &= \frac{S^E}{S_2^P} \\ &= \frac{2\lambda^*}{2\lambda^* + r} \frac{2p\Pi_2 + (1 - p)p\Pi_1}{2p\Pi_1} = \beta,\end{aligned}$$

where the last equality again uses (12) and the definition of β from (1). Since this expression applies when $\beta < 1 + p$, β can be arbitrarily small as long as p is small (to satisfy the constraint that $\beta > p$), and thus in this case $\mathfrak{s} \downarrow 0$. In both cases, the source of inefficiency of the symmetric equilibrium is because it generates insufficient incentives for experimentation. In the first case this exhibits itself as delayed experimentation, and in the second, as lack of experimentation by one of the firms.

This discussion establishes (proof in the text).⁸

Proposition 10 *1. Asymmetric equilibria are Pareto optimal and maximize social surplus when (25) holds, but fail to maximize social surplus when (25) does not hold.*

2. The unique symmetric equilibrium is always Pareto suboptimal and never maximizes social surplus. When (25) holds, this equilibrium involves delayed experimentation, and when (25) does not hold, there is insufficient experimentation.

3. When (25) holds, the relative surplus in the equilibrium compared to the surplus in the optimal allocation, \mathfrak{s} , can be as small as $1/2$. When (25) does not hold, the symmetric equilibrium can be arbitrarily inefficient. In particular, $\mathfrak{s} \downarrow 0$ as $p \downarrow 0$ and $\beta \downarrow 0$.

⁸Naturally, if we take into account consumer surpluses as discussed in footnote 7 and it is not the case that $C_2 = \kappa\Pi_2$ and $C_1 = \kappa\Pi_1$, then equilibrium social surplus relative to social surplus in the optimal allocation can be even lower because of the misalignment between firm profits and consumer surpluses resulting from different types of successful research projects.

5.2 Patents

The previous subsection established the inefficiency of the symmetric equilibrium resulting from delayed and insufficient experimentation. In this subsection, we discuss how patents can solve or ameliorate this problem. Our main argument is that a patent system provides incentives for greater experimentation or for experimentation without delay.

We model a simple patent system, whereby a patent is granted to any firm that undertakes a successful innovation. If a firm copies a patented innovation, it has to make a payment η to the holder of the patent. We discuss the relationship between this payment and licensing fees in the next subsection. An appropriately-designed patent system (i.e., the appropriate level of η) can achieve two objectives simultaneously. First, it can allow firms to copy others when it is socially beneficial for the knowledge created by innovations to spread to others (and prevent it when it is not beneficial). Second, it can provide compensation to innovators, so that incentives to free-ride on others are weakened. In particular, when staggered experimentation is optimal, a patent system can simultaneously provide incentives to one firm to innovate early and to the other firm to copy an existing innovation. When η is chosen appropriately, the patent system provides incentives for the *ex post* transfer of knowledge. However, more crucially, it also encourages innovation because an innovation that is copied becomes more profitable than copying another innovation and paying the patent fee. The key here is that the incentives provided by the patent system are “conditional” on whether the other firm has experimented or not, and thus induce an “asymmetric” response from the two firms. This makes innovation relatively more profitable when the other firm copies and less profitable when the other firm innovates. This incentive structure encourages one of the firms to be the innovator precisely when the other firm is copying. Consequently, the resulting equilibria resemble asymmetric equilibria. Moreover, these asymmetric incentives imply that, when the patent system is designed appropriately, a symmetric equilibrium no longer exists. It is less profitable for a firm to innovate when the other firm is also innovating, because innovation no longer brings patent revenues. Conversely, it is not profitable for a firm to wait when the other firm waits, because there is no innovation to copy in that case.

Our main result in this subsection formalizes these ideas. We state this result in the following proposition and then provide most of the proof, which is intuitive, in the text.

Proposition 11 *Consider the model with two firms. Suppose that Assumption 1 holds. Then:*

1. When (25) holds, a patent system with

$$\eta \in \left[\frac{(1-p)\Pi_1}{2}, \Pi_2 - p\Pi_1 \right)$$

(which is feasible in view of (25)), implements the optimal allocation, which involves staggered experimentation, in all equilibria. That is, in all equilibria one firm experiments first, and the other one copies a successful innovation and experiments immediately following an unsuccessful experimentation.

2. When (25) does not hold, then the optimal allocation, which involves simultaneous experimentation, is implemented as the unique equilibrium by a patent system with

$$\eta > \Pi_2 - p\Pi_1.$$

That is, there exists a unique equilibrium in which both firms immediately experiment.

Let us start with the first claim in Proposition 11. Observe that since $\eta < \Pi_2 - p\Pi_1$, the equilibrium involves copying of a successful innovation by a firm that has not acted yet. However, incentives for delaying to copy are weaker because copying now has an additional cost η , and innovation has an additional benefit η if the other firm is imitating. Suppose that firm $\sim i$ will innovate at some date $T > 0$ (provided that firm i has not done so until then). Then the payoffs to firm i when it chooses experimentation and waiting are

$$\begin{aligned} \text{experiment now} &= p(\Pi_2 + \eta) \\ \text{wait} &= e^{-rT} (p(\Pi_2 - \eta) + (1-p)p\Pi_1). \end{aligned}$$

It is clear that for any $T > 0$, experimenting is a strict best response, since

$$p(\Pi_2 + \eta) \geq p(\Pi_2 - \eta) + (1-p)p\Pi_1$$

since $\eta \geq \frac{(1-p)\Pi_1}{2}$. So experimenting immediately against a firm that is waiting is optimal. To show that all equilibria implement the optimal allocation, we also need to show that both firms experimenting immediately is not an equilibrium. Suppose they did so. Then the payoff to each firm, as a function of whether they experiment or wait, would be

$$\begin{aligned} \text{experiment now} &= p\Pi_1 \\ \text{wait} &= p(\Pi_2 - \eta) + (1-p)p\Pi_1. \end{aligned}$$

Waiting is a strict best response since

$$p(\Pi_2 - \eta) + (1-p)p\Pi_1 > p\Pi_1$$

which holds in view of the fact that $\eta < \Pi_2 - p\Pi_1$. This argument makes it intuitive that patents induce an equilibrium structure without delay: waiting is (strictly) optimal when the other firm is experimenting immediately and experimenting immediately is (strictly) optimal when the other firm is waiting. However, to establish this formally, we need to prove that there are no mixed strategy equilibria. This is done in the next lemma.

Lemma 2 *When equation (25) holds, there does not exist any equilibrium with mixing.*

Proof. Let us write the expected present discounted value of experimenting at time t for firm i when firm $\sim i$ experiments at the flow rate $\lambda(t)$ as in (14) in the proof of Proposition 6 except that we now take patent payments into account and use equation (25) so that copying a successful innovation is still profitable. This expression is

$$V(t) = \int_0^t \lambda(z) e^{-m(z)} e^{-rz} [p(\Pi_2 - \eta) + (1-p)p\Pi_1] dz + e^{-m(t)} e^{-rt} p(\Pi_2 + \eta),$$

where $m(t)$ is given by (13) in the proof of Proposition 6. This expression must be constant for all t in the support of the mixed-strategy equilibrium. The argument in the proof of Proposition 6 establishes that $\lambda(t)$ must satisfy

$$\lambda(t) [p(\Pi_2 - \eta) + (1-p)p\Pi_1] = (r + \lambda(t)) p(\Pi_2 + \eta).$$

It can be verified easily that if (25) holds this equation cannot be satisfied for any $\lambda(t) \in \mathbb{R}_+$ (for any t). Therefore, there does not exist any equilibrium with mixing. ■

Let us next turn to the second claim in the proposition. Suppose that (25) is not satisfied and let $\eta > \Pi_2 - p\Pi_1$. Then it is not profitable for a firm to copy a successful innovation. Therefore, both firms have a unique optimal strategy which is to experiment immediately, which coincides with the optimal allocation characterized in Proposition 9.

The preceding discussion and Proposition 11 show how an appropriately-designed patent system can be useful by providing stronger incentives for experimentation. When simultaneous experimentation by all parties is socially beneficial, a patent system can easily achieve this by making copying (or “free-riding”) unprofitable. On the other hand, when ex post transfer of knowledge is socially beneficial, the patent system can instead ensure this while also preventing delays in all equilibria. It is important to emphasize that, in the latter case, the patent system provides such incentives *selectively*, so that only one of the firms engages in experimentation and the other firm potentially benefits from the innovation of the first firm. In contrast to patents, simple subsidies to research could not achieve this objective. This is stated in the next proposition and highlights the particular utility of a patent system in this environment.

Proposition 12 *Suppose equation (25) holds. Consider a direct subsidy $w > 0$ given to a firm that experiments. There exists no $w \geq 0$ such that all equilibria with subsidies correspond to the optimal allocation.*

Proof. This is straightforward to see. If $w \geq \Pi_2 - p\Pi_1$, there exists an equilibrium in which both firms experiment immediately and if $w < \Pi_2 - p\Pi_1$, the symmetric mixed-strategy equilibrium with delayed experimentation survives. ■

It is clear that the same argument applies to subsidies to successful innovation or any combination of subsidies to innovation and experimentation.

5.3 Patents and License Fees

The analysis in the previous subsection assumed that a firm can copy a successful innovation and in return it has to make some pre-specified payment η to the original innovator. In practice patents often provide exclusive rights to the innovator, who is then allowed to license its product or discovery to other firms. With this interpretation, the payment η would need to be negotiated between the innovator and the (potential) copying firm rather than determined in advance. While licensing is an important aspect of the patent system in practice, it is not essential for the theoretical insights we would like to emphasize.

To illustrate this, let us suppose that the copying firm is developing a different but highly substitutable product to the first innovation. Suppose further that the patent system gives exclusive rights to the innovator but if the second firm copies a successful innovation, the court system needs to determine damages. How the court system functions is also part of the patent system. In particular, suppose that if a firm copies a successful innovation without licensing and the innovator brings a lawsuit, it will succeed with probability $\rho \in (0, 1)$ and the innovator will receive damages equal to $\kappa(\Pi_1 - \Pi_2)$, where $\kappa > 0$. We ignore legal fees. Given this legal environment, let us interpret η as a license fee negotiated between the potential copying firm and the innovator. For simplicity, suppose that this negotiation can be represented by a take-it-or-leave-it offer by the innovator (this has no effect on the conclusions of this subsection). If the two firms agree to licensing, their joint surplus is $2\Pi_2$. If they disagree, then the outside option of the copying firm is $\max\{p\Pi_1; \Pi_2 - \rho\kappa(\Pi_1 - \Pi_2)\}$, where the max operator takes care of the fact that the best alternative for the “copying” firm may be to experiment if there is no explicit licensing agreement. Without licensing, the innovator will receive an expected return of $\Pi_2 + \rho\kappa(\Pi_1 - \Pi_2)$ if $\Pi_2 - \rho\kappa(\Pi_1 - \Pi_2) \geq p\Pi_1$ and Π_1 otherwise. This implies that the

negotiated licensing fee, as a function of the parameters of the legal system, will be

$$\eta(\rho, \kappa) = \begin{cases} \rho\kappa(\Pi_1 - \Pi_2) & \text{if } p\Pi_1 < \Pi_2 - \rho\kappa(\Pi_1 - \Pi_2), \\ \Pi_2 - p\Pi_1 & \text{if } p\Pi_1 \geq \Pi_2 - \rho\kappa(\Pi_1 - \Pi_2) \text{ and } 2\Pi_2 > \Pi_1, \\ \infty & \text{otherwise,} \end{cases}$$

where ∞ denotes a prohibitively expensive licensing fee, such that no copying takes place. Clearly, by choosing ρ and κ , it can be ensured that $\eta(\rho, \kappa)$ is greater than $\Pi_2 - p\Pi_1$ when (25) does not hold and is between $(1 - p)\Pi_1/2$ and $\Pi_2 - p\Pi_1$ when it holds. This illustrates how an appropriately-designed legal enforcement system can ensure that equilibrium licensing fees play exactly the same role as the pre-specified patent fees did in Proposition 11.

6 Model with Heterogeneous Information

In this and the next two sections, we relax the assumption that all firms receive signals with identical precision. Instead, now signal quality differs stochastically across firms. We continue to assume that each firm receives a positive signal about a single project. But the information content of these signals differs. We parameterize signal quality by the probability with which the indicated project is successful and denote it by p (or by p_i for firm i). Throughout this and the next two sections, we assume that p is drawn from a distribution represented by the cumulative distribution function $G(p)$. We also assume that G has strictly positive and continuous density $g(p)$ over its support $[a, b] \subset [0, 1]$. The realization of p for each firm is independent of the realizations for others and is *private information*.

We continue to define $\beta \equiv \Pi_2/\Pi_1$ as in (1). We also focus on the equivalent of *symmetric equilibria* where strategies do not depend on firm identity. Asymmetric equilibria are discussed briefly in Section 8.

In this section, we will establish a *monotonicity property* showing that in equilibrium firms with higher p (stronger signals) experiment earlier (no later) than firms with lower p . This result turns out to be true under two scenarios: when there are two firms or when there are multiple firms but the upper support, b , of G is not “too large”. In the next section, we show how this monotonicity result does not generalize to an environment in which there are multiple firms and $b > \beta$.

6.1 Two Firms

The following lemma is straightforward and follows from the definition of β in (1). It will play an important role in the analysis that follows (proof omitted).

Lemma 3 *Suppose that firm $\sim i$ has innovated successfully. If $p_i > \beta$, firm i prefers to experiment with its own project. If $p_i < \beta$, firm i prefers to copy the successful project.*

As a first step towards characterizing the equilibrium with two firms, let us suppose that p has support $[a, b] \subset [0, \beta]$. Under this assumption, we will prove that there exists a symmetric equilibrium represented by a strictly decreasing function $\tau(p)$ with $\tau(b) = 0$ which maps signals to time of experimentation provided that the other player has not yet experimented. The following proposition formalizes this idea and is proved by using a series of lemmas.

Proposition 13 *Suppose that the support of G is $[a, b] \subset [0, \beta]$. Define*

$$\tau(p) = \frac{1}{r\beta} \left[\log G(b)(1-b) - \log G(p)(1-p) + \int_p^b \log G(z) dz \right]. \quad (27)$$

Then the unique symmetric equilibrium takes the following form:

1. *each firm copies a successful innovation and immediately experiments if the other firm experiments unsuccessfully;*
2. *firm i with signal quality p_i experiments at time $\tau(p_i)$ given by (27) unless firm $\sim i$ has experimented before time $\tau(p_i)$.*

Proof. The proof uses the following lemmas.

Lemma 4 *$\tau(p)$ cannot be locally constant. That is, there exists no interval $P = [\bar{p}, \bar{p} + \epsilon]$ with $\epsilon > 0$ such that $\tau(p) = t$ for all $p \in P$.*

Proof. Suppose, to obtain a contradiction, that the equilibrium involves $\tau(p) = t$ for all $p \in P$. Then, let $p_i \in P$. Firm i 's (time t) payoff after the game has reached (without experimentation) time t is

$$v(t | p_i) = p_i [(G(\bar{p} + \epsilon) - G(\bar{p})) \Pi_1 + (1 - G(\bar{p} + \epsilon) + G(\bar{p})) \Pi_2],$$

since with probability $G(\bar{p} + \epsilon) - G(\bar{p})$ firm $\sim i$ has $p \in P$ and thus also experiments at time t . In this case, firm i , when successful, is not copied and receives Π_1 . With the complementary probability, it is copied and receives Π_2 . Now consider the deviation $\tau(p_i) = t + \delta$ for $\delta > 0$ and arbitrarily small. The payoff to this is

$$v_d(t | p_i) = e^{-r\delta} [(G(\bar{p} + \epsilon) - G(\bar{p})) (\zeta \Pi_2 + (1 - \zeta) p_i \Pi_1) + (1 - G(\bar{p} + \epsilon) + G(p_i)) p_i \Pi_2],$$

where $\zeta \equiv \mathbb{E}[p | p \in P]$ is the expected probability of success of a firm with type in the set P . Since $\Pi_2 > p_i \Pi_1$, we have $\zeta \Pi_2 + (1 - \zeta) p_i \Pi_1 > p_i \Pi_1$. Moreover, by the assumption that G has strictly positive density, $G(\bar{p} + \epsilon) - G(\bar{p}) > 0$. Thus for δ sufficiently small, the deviation is profitable. This contradiction establishes the lemma. ■

Lemma 5 $\tau(p)$ is continuous in $[a, b]$.

Proof. Suppose $\tau(p)$ is discontinuous at \bar{p} . Assume without loss of generality that $\tau(\bar{p}+) \equiv \lim_{\epsilon \downarrow 0} \tau(\bar{p} + \epsilon) > \tau(\bar{p}-) \equiv \lim_{\epsilon \uparrow 0} \tau(\bar{p} - \epsilon)$. Then firms with signal $p = \bar{p} + \delta$ for sufficiently small $\delta > 0$ can experiment at time $\tau(\bar{p}-) + \epsilon$ for $\epsilon < \tau(\bar{p}+) - \tau(\bar{p}-)$ and increase their payoff since $r > 0$. ■

Lemma 6 $\tau(p)$ is strictly monotone on $[a, b]$.

Proof. Suppose, to obtain a contradiction, that there exist $q_1 > q_2$ such that $\tau(q_1) = \tau(q_2) = \bar{\tau}$. Suppose that $\sim i$ follows the equilibrium strategy characterized by $\tau(p)$ and consider firm i 's expected profit when $p_i = q$ and it chooses to experiment at time t . This can be written as

$$V(q, t) = \int_{p \in P_{before}^t} e^{-r\tau(p)} (p\Pi_2 + (1-p)q\Pi_1) dG(p) + e^{-rt} q\Pi_2 \int_{p \in P_{after}^t} dG(p), \quad (28)$$

where $P_{before}^t = \{p: \tau(p) \leq t\}$ and $P_{after}^t = \{p: \tau(p) > t\}$. Notice that $V(q, t)$ is linear in q .

For $\tau(p)$ to characterize a symmetric equilibrium strategy and given our assumption that $\tau(q_1) = \tau(q_2) = \bar{\tau}$, we have

$$V(q_1, \bar{\tau}) \geq V(q_1, t') \quad \text{and} \quad V(q_2, \bar{\tau}) \geq V(q_2, t') \quad (29)$$

for all $t' \in \mathbb{R}_+$.

Now take $q = \alpha q_1 + (1 - \alpha)q_2$ for some $\alpha \in (0, 1)$. By the linearity of $V(q, t)$, this implies that for any $t \neq \bar{\tau}$, we have

$$\begin{aligned} V(\alpha q_1 + (1 - \alpha)q_2, t) &= \alpha V(q_1, t) + (1 - \alpha) V(q_2, t) \\ &\leq \alpha V(q_1, \tau(q_1)) + (1 - \alpha) V(q_2, \tau(q_2)) \\ &= V(\alpha q_1 + (1 - \alpha)q_2, \bar{\tau}), \end{aligned}$$

where the middle inequality exploits (29). This string of inequalities implies that $\tau(\alpha q_1 + (1 - \alpha)q_2) = \bar{\tau}$ for $\alpha \in [0, 1]$. Therefore, τ must be constant between q_1 and q_2 . But this contradicts Lemma 4, establishing the current lemma. ■

The three lemmas together establish that τ is continuous and strictly monotone. This implies that τ is invertible, with inverse $\tau^{-1}(t)$. Moreover, $\tau(b) = 0$, since otherwise a firm with signal $b - \epsilon$ could experiment earlier and increase its payoff. Now consider the maximization problem of firm i with signal q . This can be written as an optimization problem where the firm in question

chooses the threshold signal $p = \tau^{-1}(t)$ rather than choosing the time of experimentation t . In particular, this maximization problem can be written as

$$\max_{p \in [a, b]} \int_p^b e^{-r\tau(p_{\sim i})} (p_{\sim i} \Pi_2 + (1 - p_{\sim i}) q \Pi_1) dG(p_{\sim i}) + e^{-r\tau(p)} G(p) q \Pi_2, \quad (30)$$

where the first term is the expected return when the firm $\sim i$ has signal quality $p_{\sim i} \in [p, b]$ and the second term is the expected return when $p_{\sim i} < p$, so that firm $\sim i$ will necessarily copy from i 's successful innovation.

Next, suppose that τ is differentiable (we will show below that τ must be differentiable). Then the objective function (30) is also differentiable and the first-order optimality condition can be written (after a slight rearrangement) as

$$r\tau'(p) = \frac{g(p)}{G(p)} \left[1 - \frac{p}{q} - (1 - p)\beta^{-1} \right].$$

In a symmetric equilibrium, the function $\tau(p)$ must be a best response to itself, which here corresponds to $p = q$. Therefore, when differentiable, $\tau(p)$ is a solution to

$$r\tau'(p) = -\frac{g(p)}{G(p)}(1 - p)\beta^{-1}. \quad (31)$$

Integrating this expression, then using integration by parts and the boundary condition $\tau(b) = 0$, we obtain the unique solution (when $\tau(p)$ is differentiable) as

$$\begin{aligned} \tau(p) &= \frac{1}{r\beta} \int_p^b (1 - z) \frac{g(z)}{G(z)} dz \\ &= \frac{1}{r\beta} \left[\log G(b)(1 - b) - \log G(p)(1 - p) + \int_p^b \log G(z) dz \right]. \end{aligned}$$

To complete the proof, we need to establish that this is the unique solution. Lemmas 5 and 6 imply that $\tau(p)$ must be continuous and strictly monotone. The result follows if we prove that $\tau(p)$ is also differentiable. Recall that a monotone function is differentiable almost everywhere, i.e., it can have at most a countable number of points of non-differentiability (see, for example, Folland, 1999, p. 101, Theorem 3.23). Take \bar{p} to be a point of non-differentiability. Then there exists some sufficiently small $\epsilon > 0$ such that $\tau(p)$ is differentiable on $(\bar{p} - \epsilon, \bar{p})$ and on $(\bar{p}, \bar{p} + \epsilon)$. Then (31) holds on both of these intervals. Integrating it over these intervals, we obtain

$$\begin{aligned} \tau(p) &= \tau(\bar{p} - \epsilon) - \frac{1}{r\beta} \int_{\bar{p} - \epsilon}^p (1 - z) \frac{g(z)}{G(z)} dz && \text{for } p \in (\bar{p} - \epsilon, \bar{p}), \text{ and} \\ \tau(p) &= \tau(\bar{p}) - \frac{1}{r\beta} \int_{\bar{p}}^p (1 - z) \frac{g(z)}{G(z)} dz && \text{for } p \in (\bar{p}, \bar{p} + \epsilon). \end{aligned}$$

Now taking the limit $\epsilon \rightarrow 0$ on both intervals, we have either (i) $\tau(\bar{p}+) \neq \tau(\bar{p}-)$; or (ii) $\tau(\bar{p}+) = \tau(\bar{p}-)$. The first of these two possibilities contradicts continuity, so (ii) must apply. But then $\tau(\bar{p})$ is given by (27) and is thus differentiable. This argument establishes that $\tau(p)$ is differentiable everywhere and proves the uniqueness of equilibrium. ■

A particularly simple example to Proposition 13 is obtained when G is uniform over $[a, b]$ for $0 < a < b \leq \beta$. In that case

$$\tau(p) = \frac{1}{r\beta} [p - \log p - b + \log b] \text{ for all } p \in [a, b]. \quad (32)$$

An interesting feature of symmetric equilibria in this case is evident from (32): for a arbitrarily close to 0, experimentation may be delayed for arbitrarily long time. It can be verified from (27) that this is a general feature (for types arbitrarily close to the lower support a , $-\log G(p)$ is arbitrarily large).

Proposition 13 characterizes the unique equilibrium when both firms have $p \leq \beta$. The complete characterization of equilibrium is provided in the next proposition. An important implication of this equilibrium characterization is the monotonicity property of symmetric equilibria mentioned in the Introduction (we will see in Section 8 that this is typically not the property of asymmetric equilibria). Firms with higher signal quality never experiment after firms with lower signal quality.

Proposition 14 *Let the support of G be $[a, b] \subset [0, 1]$ and define $\bar{b} \equiv \min\{\beta, b\}$ and*

$$\bar{\tau}(p) \equiv \frac{1}{r\beta G(\bar{b})} \left[\log G(\bar{b}) (1 - \bar{b}) - \log G(p) (1 - p) + \int_p^{\bar{b}} \log G(z) dz \right]. \quad (33)$$

Then the unique symmetric equilibrium involves:

$$\tau(p) = \begin{cases} 0 & \text{if } p \geq \beta \\ \bar{\tau}(p) & \text{if } p \in [a, \beta) \end{cases} .$$

That is, firms with $p \geq \beta$ experiment immediately and firms with $p \in [a, \beta)$ experiment at time $\bar{\tau}(p)$ unless there has been an experimentation at $t < \bar{\tau}(p)$. If there is experimentation at $t < \bar{\tau}(p)$, then a firm with $p \in [a, \beta)$ copies it if the previous attempt was successful and experiments immediately if it was unsuccessful.

Proof. The proof consists of three steps.

The first step of the proof is to show that firms with $p \geq \beta$ will always experiment before firms with $p \in [a, \beta)$. The proof is by a single-crossing argument. First, recall that the value of experimenting at time t for a firm with $p \in [a, \beta)$ is given by (28). Defining $P_{after \wedge \beta+}^t = \{p:$

$\tau(p) > t$ and $p \geq \beta$ and $P_{after\wedge\beta-}^t = \{p: \tau(p) > t \text{ and } p < \beta\}$, the value of experimenting for a firm with $q \in [a, \beta)$ can be rewritten as

$$V(q, t) = q\Pi_2 \left\{ \int_{p \in P_{before}^t} e^{-r\tau(p)} \left(\frac{p}{q} + \frac{1-p}{\beta} \right) dG(p) + e^{-rt} \left[\frac{1}{\beta} \int_{p \in P_{after\wedge\beta+}^t} dG(p) + \int_{p \in P_{after\wedge\beta-}^t} dG(p) \right] \right\}, \quad (34)$$

which exploits the fact that when $p \in P_{before}^t$ or when $p \in P_{after\wedge\beta+}^t$, there will be no copying, and when $p \in P_{after\wedge\beta-}^t$, the innovation (which takes place again with probability q) will be copied, for a payoff of $\Pi_2 = \beta\Pi_1$.

Next, turning to firms with $p = q' \geq \beta$, first recall that these firms prefer not to copy prior successful experimentation (from Lemma 3). Therefore, their corresponding value can be written as

$$\tilde{V}(q', t) = q'\Pi_2 \left\{ \frac{1}{\beta} \int_{p \in P_{before}^t} e^{-r\tau(p)} dG(p) + e^{-rt} \left[\frac{1}{\beta} \int_{p \in P_{after\wedge\beta+}^t} dG(p) + \int_{p \in P_{after\wedge\beta-}^t} dG(p) \right] \right\}, \quad (35)$$

Note also that when the experimentation time is reduced, say from t to $t' < t$, the first integral gives us the cost of such a change and the second expression (e^{-rt} times the square bracketed term) gives the gain. Now the comparison of (34) to (35) establishes the single-crossing property, meaning that at any t a reduction to $t' < t$ is always strictly more valuable for $q' \geq \beta$ than for $q \in [a, \beta)$. First, the gains, given by the expression in (34) and (35) are identical. Second, the term in parenthesis in the first integral in (34) is a convex combination of $1/q > 1/\beta$ and $1/\beta$, and thus is strictly greater than $1/\beta$, so that the cost is always strictly greater for $q \in [a, \beta)$ than for $q' \geq \beta$. From this strict single-crossing argument it follows that there exists some T such that $\tau(p) \leq T$ for all $p \geq \beta$ and $\tau(p) > T$ for all $p \in [a, \beta)$.

The next step of the proof establishes that all firms with $p \geq \beta$ will experiment immediately, that is, $\tau(p) = 0$ for all $p \geq \beta$. To show this, first note that all terms in (35) are multiplied by $q' \geq \beta$, so the optimal set of solutions for any firm with $p \geq \beta$ must be identical. Moreover, since $\tau(p) > T$ for all $p \in [a, \beta)$, $P_{after\wedge\beta-}^t$ is identical for all $t \in [0, T]$, and $t > 0$ is costly because $r > 0$. Therefore, the unique optimal strategy for all $p \geq \beta$ is to experiment immediately. Therefore, $\tau(p) = 0$ for all $p \geq \beta$.

The final step is to combine the equilibrium behavior of firms with $p \geq \beta$ with those of $p \in [a, \beta)$. First suppose that $b \leq \beta$. Then the characterization in Proposition 13 applies exactly. Next suppose that $b > \beta$, so that some firms might have signals $p \geq \beta$. The previous step of the proof has established that these firms will experiment immediately. Subsequently, firms with $p \in [a, \beta)$ will copy a successful innovation at time $t = 0$ or experiment if there is an unsuccessful experimentation at $t = 0$. If there is no experimentation at $t = 0$, then equilibrium

behavior (of firms with $p \in [a, \beta)$) is given by Proposition 13 except that the upper support is now β and the relevant distribution is $G(p)$ conditional on $p \in [a, \beta)$, thus all terms are divided by $G(\beta)$. This completes the proof of the proposition. ■

6.2 Multiple Firms

The previous subsection established that when there are two firms, the equilibrium satisfies the monotonicity property whereby firms with stronger signals always experiment earlier than those with weaker signals. The same result generalizes to the case where there are $N > 2$ firms as long as the support of $G(p)$ does not include “very strong” signals. In particular, let

$$\beta_N \equiv \frac{\Pi_N}{\Pi_1}. \quad (36)$$

Monotonicity then requires that the support of G is $[a, b] \subset [0, \beta_N]$. In the next section, we will see that monotonicity no longer holds when $[a, b] \not\subset [0, \beta_N]$.

Proposition 15 *Suppose that there are $N \geq 2$ firms and the support of G is $[a, b] \subset [0, \beta_N]$, where β_N is defined in (36). Then there exists a unique symmetric equilibrium. This equilibrium takes the following form:*

- *a firm with signal p experiments at time $\tau_n(p)$ if there has been $N - n$ unsuccessful experiments before, where $\tau_n(p)$ is strictly decreasing and continuous for $n = 0, 1, \dots, N - 2$.*
- *all firms copy immediately if there is a successful innovation.*

Proof. The proof is similar to that of Proposition 13. It involves establishing the equivalents of Lemmas 4-6 in the continuation game in which there have been $n = 0, 1, \dots, N - 2$ unsuccessful experiments for each. Then we solve for $\tau_2(p)$ using an identical argument to that of Proposition 13. Given $\tau_2(p)$, we can then solve recursively for $\tau_n(p)$ for each $n = 3, \dots, N - 2$. The details are identical to the arguments in Proposition 13 and are omitted. ■

7 Nonmonotonicity

Equilibria characterized in the previous section involve inefficient delay as in the model with firms with symmetric signals. Nevertheless, there is a monotonicity property, whereby firms with stronger signals experiment earlier than firms with weaker signals. In this section, we show that monotonicity does not always apply. In particular, when there are $N > 2$ firms and the support of G satisfies $[a, b] \not\subset [0, \beta]$, any symmetric equilibrium *necessarily* involves nonmonotonicity with positive probability.

The main result of this section is provided in the next proposition, which establishes that monotonicity no longer holds in this case.

Proposition 16 *Suppose that there are $N \geq 3$ firms and the support of G satisfies $[a, b] \not\subset [0, \beta]$. Then:*

1. *There does not exist a symmetric equilibrium in which all firms with $p \geq \beta$ experiment at $t = 0$.*
2. *In any symmetric equilibrium, there is positive probability that a firm with $p > \beta$ will experiment earlier than a firm with $p' > p$.*

Proof. (Part 1) Suppose that $N = 3$, and that $\Pi_2 = \Pi_3$. Suppose, to obtain a contradiction, that there exists a symmetric equilibrium where all firms with $p \geq \beta$ experiment at $t = 0$. Consider firm i with $p_i > \beta$. Let χ_0 be the probability that none of the other two firms have $p \geq \beta$, χ_1 be the probability that one of the other two firms has $p \geq \beta$ and χ_2 be the probability that both firms have $p \geq \beta$. Let us also define $\zeta = \mathbb{E}[p \mid p \geq \beta]$. Since $p_i > \beta$, by hypothesis, firm i experiments at time $t = 0$. Its expected payoff is

$$V(p_i, 0) = \chi_0 p_i \Pi_2 + \chi_1 p_i \left[\zeta \left(\frac{\Pi_1}{2} + \frac{\Pi_2}{2} \right) + (1 - \zeta) \Pi_2 \right] + \chi_2 p_i \Pi_1.$$

Intuitively, when none of the other two firms have $p \geq \beta$, when successful, the firm is copied immediately, receiving payoff Π_2 . When both of the other two firms have $p \geq \beta$, there is no copying, so when successful, firm i receives Π_1 . When one of the other two firms has $p \geq \beta$, then this other firm also experiments at time $t = 0$ and is successful with probability $\zeta = \mathbb{E}[p \mid p \geq \beta]$. In that case, in a symmetric equilibrium the third firm copies each one of the two successful innovations with probability $1/2$. With the complementary probability, $1 - \zeta$, the other firm with $p \geq \beta$ is unsuccessful, and the third firm necessarily copies firm i .

Now consider the deviation to wait a short interval $\epsilon > 0$ before innovation. This will have payoff

$$\begin{aligned} \lim_{\epsilon \downarrow 0} V(p_i, \epsilon) &= \chi_0 p_i \Pi_2 + \chi_1 p_i (\zeta \Pi_1 + (1 - \zeta) \Pi_2) + \chi_2 p_i \Pi_1 \\ &> V(p_i, 0). \end{aligned}$$

The first line of the previous expression follows since, with this deviation, when there is one other firm with $p \geq \beta$, the third firm necessarily will copy the first innovator. The inequality follows since $\Pi_1 > \Pi_2$, establishing that there cannot be an equilibrium in which all firms with

$p \geq \beta$ experiment at time $t = 0$. This argument generalizes, with a little modification, to cases in which $N > 3$ and Π_n s differ.

(Part 2) Part 1 implies that any symmetric equilibrium must involve mixing over the time of experimentation by firms with $p \geq \beta$. By the same argument following from (35) in the proof of Proposition 14, all firms with $p \geq \beta$ will have the same set of experimentation times that maximize their expected payoffs, so must use the same mixed strategies in any symmetric equilibrium. Therefore, there is a positive probability that any one of them will experiment first, establishing the claim in Part 2. ■

The next proposition characterizes the form of the mixed strategy equilibrium when the monotonicity property no longer holds. To simplify the exposition, we focus on an economy in which $N = 3$ and G is uniform on $[0, 1]$. The characterization result in this proposition can be (relatively straightforwardly) extended to $N > 3$. We also conjecture that it can be extended to any distribution G , though this is less trivial.

Proposition 17 *Consider an economy with $N = 3$ firms, $\Pi_2 = \Pi_3$ and G uniform over $[0, 1]$. Then, the following characterizes the unique symmetric equilibrium.*

- *Firms with $p \geq \beta$ experiment at the flow rate $\xi(t)$ as long as no other firm has experimented until t . They experiment immediately following another (successful or unsuccessful) experiment. There exists $T < \infty$ such that*

$$e^{-\int_0^T \xi(t)dt} = 0.$$

That is, all firms with $p \geq \beta$ will have necessarily experimented within the interval $[0, T]$ (or equivalently, $\lim_{t \rightarrow T} \xi(t) = +\infty$).

- *Firms with $p < \beta$ immediately copy a successful innovation and experiment according to a flow rate $\lambda_2(t)$ following an unsuccessful experimentation and at the rate $\lambda_3(t)$ if there has been no experimentation until time T .*

Proof. From Proposition 16, firms with $p \geq \beta$ must mix over experimentation times on some subset of \mathbb{R}_+ . Let us define $\mu(t)$ as the probability that firm $\sim i$ that has not experimented until time t has $p_{\sim i} \geq \beta$. The assumption that G is uniform over $[0, 1]$ implies that $\mu(0) = 1 - \beta$.

Now consider the problem of firm i with $p_i \geq \beta$. If there has yet been no experimentation and this firm experiments at time t , its payoff (discounted to time $t = 0$) is

$$V(p_i, t) = p_i e^{-rt} [\Pi_1 \mu(t)^2 + \Pi_2 (1 - \mu(t)^2)],$$

since $\mu(t)^2$ is the probability with which both other firms have $p \geq \beta$ and will thus not copy. With the complementary probability, its innovation will be copied. Alternately, if it delays experimentation by some small amount $dt > 0$, then its payoff is:

$$V(p_i, t + dt) = p_i e^{-r(t+dt)} \left[2\bar{p}\xi(t)\mu(t)dt\Pi_1 + (1 - 2\xi(t)\mu(t)dt)[\Pi_1\mu(t+dt)^2 + \Pi_2(1 - \mu(t+dt))^2] \right. \\ \left. + 2\xi(t)\mu(t)(1 - \bar{p})dt[\Pi_1\mu(t+dt) + \Pi_2(1 - \mu(t+dt))] \right],$$

where $\bar{p} \equiv \mathbb{E}[p|p > \beta]$ and we use the fact that other firms with $p \geq \beta$ experiment at the rate $\xi(t)$. In a mixed-strategy equilibrium, these two expressions must be equal (as $dt \rightarrow 0$). Setting these equal and rearranging, we obtain

$$\frac{d(\mu^2(t))}{dt} + 2\xi(t)\mu(t)(1 - \mu(t))[\bar{p} + \mu(t)] - r\mu^2(t) = \frac{\Pi_2 r}{\Pi_1 - \Pi_2}. \quad (37)$$

In addition, the evolution of beliefs $\mu(t)$ given the uniform distribution and flow rate of experimentation at $\xi(t)$ can be obtained as

$$\mu(t) = \frac{e^{-\int_0^t \xi(\tau) d\tau} (1 - \beta)}{e^{-\int_0^t \xi(\tau) d\tau} (1 - \beta) + \beta}. \quad (38)$$

Now let us define

$$f(t) \equiv e^{-\int_0^t \xi(\tau) d\tau} (1 - \beta). \quad (39)$$

Using (39), (38) can be rewritten as

$$\mu(t) = \frac{f(t)}{f(t) + \beta},$$

which in turn implies

$$f(t) = \frac{\beta\mu(t)}{1 - \mu(t)}.$$

Moreover (39) also implies that

$$\xi(t) = -\frac{f'(t)}{f(t)} = -\frac{\mu'(t)}{(1 - \mu(t))\mu(t)}. \quad (40)$$

Substituting these into (37), we obtain the following differential equation for the evolution of $\mu(t)$:

$$2\mu(t)\mu'(t) + 2\xi(t)\mu(t)(1 - \mu(t))[\bar{p} + \mu(t)] - r\mu^2(t) = \frac{r}{\beta^{-1} - 1}.$$

Further substituting $\xi(t)$ from (40), we obtain

$$\mu'(t) = -\frac{r}{2\bar{p}(\beta^{-1} - 1)}(1 + (\beta^{-1} - 1)\mu^2(t)).$$

This differential equation satisfies the Lipschitz condition and therefore it has a unique solution, which takes the form

$$\mu(t) = \frac{1}{\sqrt{\beta^{-1} - 1}} \tan \left(\sqrt{\beta^{-1} - 1} \left[-\frac{r}{2\bar{p}(\beta^{-1} - 1)} t + \frac{\arctan(\mu(0)\sqrt{\beta^{-1} - 1})}{\sqrt{\beta^{-1} - 1}} \right] \right),$$

with boundary condition $\mu(0) = 1 - \beta$. Given this solution, the flow rate of experimentation for firms with $p \geq \beta$, $\xi(t)$, is obtained from (40) as

$$\xi(t) = c_1 c_2 (1 + \tan(c_1(-c_2 t + c_3))^2) \left[-\frac{1}{-c_1 + \tan(c_1(-c_2 t + c_3))} + \frac{1}{\tan(c_1(-c_2 t + c_3))} \right],$$

where

$$c_1 \equiv \sqrt{\beta^{-1} - 1}, \quad c_2 \equiv -\frac{r}{2\bar{p}(\beta^{-1} - 1)}, \quad \text{and} \quad c_3 \equiv \frac{\arctan(\mu(0)\sqrt{\beta^{-1} - 1})}{\sqrt{\beta^{-1} - 1}}.$$

It can then be verified that

$$\lim_{t \rightarrow T} \xi(t) = \infty,$$

where $T = c_3/c_2$. It can also be verified that for all $t \in [0, T]$, where firms with $p \geq \beta$ are experimenting at positive flow rates, firms with $p < \beta$ strictly prefer to wait. The equilibrium behavior of these firms after an unsuccessful experimentation or after time T is reached is given by an analysis analogous to Proposition 14 and again involves mixing. Combining these observations gives the form of the equilibrium described in the proposition. ■

Notice also that Proposition 17 implies that nonmonotonicity in this case affects the allocation of resources mainly by inducing delay (only firms with $p \geq \beta$ act during the time interval $[0, T]$), because following a single unsuccessful experimentation in $[0, T]$, other firms with $p \geq \beta$ experiment immediately and these firms would have preferred not to copy a successful innovation.

8 Patents, Heterogeneity, and Nonmonotonicity

In this section, we discuss the optimal allocation with heterogeneity, its comparison to equilibrium, and the role of patents in the economy with heterogeneity.

8.1 Welfare

Consider a social planner that is interested in maximizing total surplus (as in Section 5). What the social planner can achieve will depend on her information and on the set of instruments that she has access to. For example, if the social planner observes the signal quality, p , for each firm, then she can achieve a much better allocation than the equilibrium characterized above.

However, it is more plausible to limit the social planner to the same information structure. In that case, the social planner will have to choose either the same equilibrium allocation as in the symmetric equilibria characterized in the previous two sections, or she will implement an asymmetric equilibrium, where one of the firms is instructed to experiment first regardless of its p (this cannot be conditioned on p since p is private information).⁹

More specifically, let us focus on the economy with two firms and suppose that the support of G is $[a, b] \subset [0, \beta]$. In this case, without eliciting information about the realization of firm types, the p 's, the planner has three strategies.

1. *Staggered asymmetric experimentation*: in this case, the social planner would instruct one of the firms to experiment immediately and then have the other firm copy if there is a successful innovation. Since the social planner does not know the p 's, she has to pick the experimenting firm randomly. We denote the social surplus generated by this strategy by S_1^P .
2. *Staggered equilibrium experimentation*: alternatively, the social planner could let the firms play the symmetric equilibrium of the previous two sections, whereby a firm of type p will experiment at time $\tau(p)$ unless there has previously been an experimentation by the other firm. We denote the social surplus generated by this strategy by S^E , since this is the same as the equilibrium outcome.¹⁰
3. *Simultaneous experimentation*: in this case, the social planner would instruct both firms to experiment immediately. We denote the social surplus generated by this strategy by S_2^P .

The social surpluses from these different strategies are given as follows. In the case of staggered asymmetric experimentation, we have

$$S_1^P = \int_a^b \left[p_1 2\Pi_2 + (1 - p_1) \left(\int_a^b p_2 dG(p_2) \right) \Pi_1 \right] dG(p_1).$$

⁹Yet another alternative is to specify exactly the instruments available to the planner and characterize the solution to a mechanism design problem by the planner. However, if these instruments allow messages and include payments conditional on messages, the planner can easily elicit the necessary information from the firms.

¹⁰Without eliciting information about firm types and without using additional instruments, the social planner cannot implement another monotone staggered experimentation allocation. For example, she could announce that if there is no innovation until some time $t > 0$, one of the firms will be randomly forced to experiment. But such schemes will not preserve monotonicity, since at time t , it may be the firm with lower p that may be picked for experimentation. In the next subsection, we discuss how she can implement better allocations using patent payments.

In contrast, the expected surplus from the unique (mixed-strategy) symmetric equilibrium can be written as

$$S^E = \int_a^b e^{-r\tau(\max\{p_1, p_2\})} \left[\max\{p_1, p_2\} 2\Pi_2 + (1 - \max\{p_1, p_2\}) \int_a^b \min\{p_1, p_2\} \Pi_1 \right] dG(p_1) dG(p_2).$$

Intuitively, this expression follows by observing that in the equilibrium as specified in Proposition 13, the firm with the stronger signal (higher p) will experiment first, so there will be delay until $\max\{p_1, p_2\}$. At that point, this firm will succeed with probability $\max\{p_1, p_2\}$, in which case the second firm will copy. If the first firm fails (probability $1 - \max\{p_1, p_2\}$), then the second firm experiments and succeeds with probability $\min\{p_1, p_2\}$. Since both p_1 and p_2 are randomly drawn independently from G , we integrate over G twice to find the expected surplus.

The surplus from simultaneous experimentation, on the other hand, takes a simple form and is given by

$$S_2^P = 2\Pi_1 \int_a^b p dG(p),$$

since in this case each firm is successful and generates payoff Π_1 with probability p distributed with distribution function G .

In this case, there is no longer any guarantee that $\max\{S_1^P, S_2^P\} > S^E$. Therefore, the symmetric equilibrium may generate a higher expected surplus (relative to allocations in which the social planner does not have additional instruments). To illustrate this, let us consider a specific example, where p has a uniform distribution over $[0, \beta]$. In this case, staggered asymmetric experimentation gives

$$S_1^P = \int_0^\beta \int_0^\beta 2p_1\Pi_2 + (1 - p_1)p_2\Pi_1 dp_2 dp_1 = \Pi_2 \left(\frac{1}{2} + \frac{3}{4}\beta \right),$$

whereas simultaneous experimentation gives

$$S_2^P = \int_0^\beta 2p\Pi_1 dp = \beta\Pi_1 = \Pi_2.$$

Comparing simultaneous experimentation and staggered asymmetric experimentation, we can conclude that $S_1^P > S_2^P$ whenever $\beta > 2/3$ and $S_1^P < S_2^P$ whenever $\beta < 2/3$, showing that, as in the case with common signals, either simultaneous or staggered experimentation might be optimal. Next, we can also compare these surpluses to S^E . Since p is uniformly distributed in $[0, \beta]$, (27) implies that

$$\tau(p) = \frac{1}{r\beta} [p - \log p - \beta + \log \beta].$$

As a consequence, $\max\{p_1, p_2\}$ has a Beta(2,1) distribution (over $[0, \beta]$) while $\min\{p_1, p_2\}$ is distributed Beta(1,2). Then evaluating the expression for S^E , we find that when $0 \leq \beta \leq 2/3$,

$S_2^P > S^E$, so simultaneous experimentation gives the highest social surplus. When $2/3 \leq \beta \leq \beta^* \simeq 0.895$, $S_1^P > S^E$, so that staggered asymmetric experimentation gives the highest social surplus. Finally, when $\beta^* \leq \beta \leq 1$, $S^E > S_1^P > S_2^P$, so the symmetric equilibrium gives higher social surplus than both staggered asymmetric experimentation and simultaneous experimentation.

Finally, it is also straightforward to see that by choosing G to be highly concentrated around a particular value \bar{p} , we can repeat the same argument as in subsection 5.1 and show that the symmetric equilibrium can be arbitrarily inefficient relative to the optimal allocation.

8.2 Equilibrium with Patents

Equilibria with patents are also richer in the presence of heterogeneity. Let us again focus on the case in which there are two firms. Suppose that there is a patent system identical to the one discussed in subsection 5.2, whereby a firm that copies a successful innovation pays η to the innovator. Let us define

$$p^\eta \equiv \frac{\Pi_2 - \eta}{\Pi_1}. \quad (41)$$

It is clear, with a reasoning similar to Lemma 3, that only firms with $p < p^\eta$ will copy when the patent system specifies a payment of η . The next proposition characterizes the structure of equilibria with patents.

Proposition 18 *Suppose that there are two firms and the patent system specifies a payment $\eta > 0$ for copying. Let p^η be given by (41), the support of G be $[a, b] \subset [0, 1]$, and define $\bar{b} \equiv \min\{b, p^\eta\}$ and*

$$\bar{\tau}^\eta(p) \equiv \frac{1}{r(\Pi_2 + \eta)G(\bar{b})} \left[\log G(\bar{b}) (\Pi_1 - 2\eta - \bar{b}\Pi_1) - \log G(p) (\Pi_1 - 2\eta - p\Pi_1) + \Pi_1 \int_p^{\bar{b}} \log G(z) dz \right]. \quad (42)$$

Then the unique symmetric equilibrium involves:

$$\tau^\eta(p) = \begin{cases} 0 & \text{if } p \geq p^\eta \\ \bar{\tau}^\eta(p) & \text{if } p \in [a, p^\eta) \end{cases}.$$

That is, firms with $p \geq p^\eta$ experiment immediately and firms with $p \in [a, p^\eta)$ experiment at time $\bar{\tau}^\eta(p)$ unless there has been an experimentation at $t < \bar{\tau}^\eta(p)$.

Moreover, a higher η tends to reduce delay. In particular:

- *for any $\eta' > \eta$ such that $b < p^\eta$ and $b < p^{\eta'}$, we have $\tau^{\eta'}(p) \leq \tau^\eta(p)$ for all $p \in [a, b]$, with strict inequality whenever $\tau^\eta(p) > 0$;*

- for any η' such that $b > p^{\eta'}$, there exists $p^*(\eta') \in [0, p^{\eta'})$ such that $\tau^\eta(p)$ is decreasing in η starting at $\eta = \eta'$ for all $p \in [p^*(\eta'), p^{\eta'})$, with strict inequality whenever $\tau^{\eta'}(p) > 0$.

Proof. The proof mimics that of Proposition 14, with the only difference that the maximization problem of firm i , with signal $p_i = q$, is now modified from (30) to

$$\max_{p \in [a, b]} \int_p^b e^{-r\tau(p_{\sim i})} (p_{\sim i} (\Pi_2 - \eta) + (1 - p_{\sim i}) q \Pi_1) dG(p_{\sim i}) + e^{r\tau(p)} G(p) q (\Pi_2 + \eta),$$

which takes into account that copying has cost η and if firm i is the first innovator, then it will be copied and will receive η . Repeating the same argument as in Proposition 14 establishes that the unique equilibrium is given by (42).

To prove the second part of the proposition, first suppose that $b < p^\eta$ and $b < p^{\eta'}$ so that $\bar{b} = b$ in both cases. Recall also that $\tau^\eta(p) = \bar{\tau}^\eta(p) > 0$ for $p \in [a, p^\eta)$. p^η is decreasing in η , so that $\tau^\eta(p) = 0$ implies that $\tau^{\eta'}(p) = 0$ for any $\eta' > \eta$. We therefore only need to show that $\bar{\tau}^\eta(p)$ is strictly decreasing for all $p \in [a, p^\eta)$. Since $\bar{\tau}^\eta(p)$ is differentiable, it is sufficient to show that its derivative with respect to η is negative. This follows by differentiating (42) (with $\bar{b} = b$). In particular,

$$\frac{d\tau^\eta(p)}{d\eta} = -\frac{1}{(\Pi_2 + \eta)} \left[\bar{\tau}^\eta(p) + \frac{2}{rG(\bar{b})} (\log G(\bar{b}) - \log G(p)) \right] < 0,$$

since $\log G(\bar{b}) > \log G(p)$ and $\bar{\tau}^\eta(p) > 0$.

Next, suppose that $b > p^{\eta'}$. In that case $\bar{b} = p^{\eta'}$ and $d\tau^\eta(p)/d\eta$ (in the neighborhood of η') will include additional terms because of the effect of η on \bar{b} . In particular:

$$\begin{aligned} \frac{d\tau^{\eta'}(p)}{d\eta} &= -\frac{1}{(\Pi_2 + \eta')} \left[\bar{\tau}^{\eta'}(p) + \frac{2}{rG(p^{\eta'})} (\log G(p^{\eta'}) - \log G(p)) \right] \\ &\quad - \frac{g(p^{\eta'})}{G(p^{\eta'})\Pi_1} \frac{\Pi_1 - 2\eta' - p^{\eta'}\Pi_1}{r(\Pi_2 + \eta')G(p^{\eta'})} + \frac{g(p^{\eta'})}{G(p^{\eta'})\Pi_1} \bar{\tau}^{\eta'}(p). \end{aligned} \quad (43)$$

The first line is again strictly negative and so is the first expression in the second line. The second expression in the second line could be positive, however. For given η' , this term is decreasing in p and tends to 0 as p approaches $p^{\eta'}$ (from (42)). Therefore, there exists $p^*(\eta')$ such that for $p \geq p^*(\eta')$, it is no larger than the first term in the second line. This establishes that for $p \in [p^*(\eta'), p^{\eta'})$, $\tau^{\eta'}(p)$ is again decreasing in η , completing the proof. ■

An important role of patents in experimentation is highlighted by this result. When η increases, $\tau(p)$ tends to become “steeper” so that there is less delay and thus “time runs faster”. In particular, whenever $p^\eta < b$, $\tau(p)$ is reduced by an increase in patent payments. When $p^\eta > b$, this does not necessarily apply for very low p 's, but is still true for high p 's. Overall, this

result implies that as in the case with common p 's, patents tend to increase experimentation incentives and reduce delay. In the limit, when η becomes arbitrarily large, the equilibrium involves simultaneous experimentation. Nevertheless, as discussed in subsection 8.1, simultaneous experimentation may not be optimal in this case.

Alternatively (and differently from Proposition 18), a patent system can also be chosen such that the socially beneficial ex post transfer of knowledge takes place. In particular, suppose that there has been an innovation and the second firm has probability of success equal to p . In this case, social surplus is equal to $2\Pi_2$ if there is copying, and it is equal to $\Pi_1 + p\Pi_1$ if the second firm is forced to experiment. This implies that to maximize ex post social welfare, firms with $p \leq 2\beta - 1$ should be allowed to copy, whereas firms with $p > 2\beta - 1$ should be induced to experiment. Clearly, from (41) choosing $\eta = \Pi_1 - \Pi_2$ achieves this. Naturally, from Proposition 18, this will typically lead to an equilibrium with staggered experimentation. This argument establishes the following proposition (proof in the text).

Proposition 19 *A patent system with $\eta = \Pi_1 - \Pi_2$ induces the socially efficient copying and experimentation behavior for all $p \in [a, b]$, but typically induces delayed experimentation.*

The juxtaposition of Propositions 18 and 19 implies that when signal quality is heterogeneous and private information, the patent system can ensure either rapid experimentation or the socially beneficial ex post transfer of knowledge (and experimentation by the right types), but will not typically be able to achieve both objectives simultaneously.

8.3 Patents and Nonmonotonicity

Our last result shows that patent systems cannot prevent the nonmonotonicity identified in the previous section.

Proposition 20 *Consider the economy with $N \geq 3$ firms and the support of G given by $[a, b] \not\subset [0, \beta]$, so that the symmetric equilibrium without patents involves nonmonotonicity. Then for any patent η , there exists an equilibrium either with no copying or nonmonotonicity. That is, there exists no patent system such that all equilibria will avoid nonmonotonicity while ensuring efficient transfer of knowledge.*

Proof. Clearly, any patent with $\eta > \Pi_1 - \Pi_2$ will prevent efficient transfer of knowledge. So suppose that $\eta \leq \Pi_1 - \Pi_2$. Then let us adopt the same definitions of ζ , χ_0 , χ_1 and χ_2 as in the proof of Proposition 16. Suppose, to obtain a contradiction, that there exists an equilibrium in

which all firms with $p \geq \beta$ experiment immediately. Then the payoff to firm $p_i \geq \beta$ when it does so is

$$V(p_i, 0) = \chi_0 p_i (\Pi_2 + 2\eta) + \chi_1 p_i \left[\zeta \left(\frac{\Pi_1}{2} + \frac{\Pi_2 + \eta}{2} \right) + (1 - \zeta) (\Pi_2 + \eta) \right] + \chi_2 p_i \Pi_1,$$

which takes into account the patent payments from the other firms into account.

Now consider a deviation for firm i that involves waiting for a short interval $\epsilon > 0$ before innovation. This will have payoff

$$\begin{aligned} \lim_{\epsilon \downarrow 0} V(p_i, \epsilon) &= \chi_0 p_i (\Pi_2 + 2\eta) + \chi_1 p_i (\zeta \Pi_1 + (1 - \zeta) (\Pi_2 + \eta)) + \chi_2 p_i \Pi_1 \\ &> V(p_i, 0), \end{aligned}$$

since $\eta < \Pi_1 - \Pi_2$. Therefore, whenever $\eta < \Pi_1 - \Pi_2$, the symmetric equilibrium will involve delay by firms with $p \geq \beta$ and thus potential nonmonotonicity. ■

9 Conclusion

This paper studied a simple model of experimentation and innovation. Each firm receives a private signal on the success probability of one of many potential research projects and decides when and which project to implement. A successful innovation can be copied by other firms. We show that, both in discrete and continuous time, symmetric equilibria, where actions do not depend on the identity of the firm, necessarily involve delayed and staggered experimentation. When the signal quality is the same for all players, the equilibrium is in mixed strategies (pure-strategy symmetric equilibria do not exist). When signal quality differs across firms, the equilibrium is represented by a function $\tau(p)$ which specifies the time at which a firm with signal quality p experiments. As in the environment with common signal quality, the equilibrium may involve arbitrarily long delays.

We also show that the social cost of insufficient experimentation incentives can be arbitrarily large. The optimal allocation may require simultaneous rather than staggered experimentation. In this case, the efficiency gap between the optimal allocation and the equilibrium can be arbitrarily large. Instead, when the optimal allocation also calls for staggered experimentation, the equilibrium is inefficient because of delays. We show that in this case the ratio of social surplus in the equilibrium to that in the optimal allocation can be as low as 1/2.

One of the main arguments of the paper is that appropriately-designed patent systems encourage experimentation and reduce delays without preventing efficient *ex post* transfer of knowledge across firms. Consequently, when signal quality is the same for all firms, an appropriately-designed patent system can ensure that the optimal allocation results in all equilibria. Patents

are particularly well-suited to providing the correct incentives when the optimal allocation also requires staggered experimentation. In this case, patents can simultaneously encourage one of the firms to play the role of a leader in experimentation, while providing incentives to others to copy successful innovations. Technically, appropriately-designed patents destroy symmetric equilibria, which are the natural equilibria in the absence of patents but may involve a high degree of inefficiency. That patents are an attractive instrument in this environment can also be seen from our result that, while patents can implement the optimal allocation, there exists no simple subsidy (to experimentation, research, or innovation) that can achieve the same policy objective.

When signal quality differs across firms, an additional dimension of efficiency is the sequence in which firms with different signals experiment. We show that when there are only two firms or when there is an arbitrary number of firms but the support of the signal quality distribution does not include “very strong” signals, the equilibrium has the following *monotonicity property*: firms with stronger signals always experiment earlier (no later) than firms with weaker signals. However, this result no longer holds when there are more than two firms and “very strong” signals are possible. In this case, we show that the equilibrium necessarily involves some amount of nonmonotonicity. Patents are again useful in encouraging experimentation in this case and reduce delays. Nevertheless, interestingly, when the equilibrium without patents involves nonmonotonicity, it is impossible for patents to both ensure efficient *ex post* transfer of knowledge and restore monotonicity.

We believe that the role of patents in encouraging socially beneficial experimentation is more general than the simple model used in this paper. In particular, throughout the paper we ignored the consumer side. It is possible that new innovations create benefits to consumers that are disproportionately greater than the use of existing successful innovations (as compared to the relative profitabilities of the same activities). In this case, the social benefits of experimentation are even greater and patents can also be useful in preventing copying of previous successful innovations. The investigation of the welfare and policy consequences of pursuing successful lines versus experimenting with new, untried research lines is an interesting and underresearched area.

References

Arrow, Kenneth (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.

Bolton, Patrick and Christopher Harris (1999) "Strategic Experimentation," *Econometrica*, 67, 347-374.

Bolton, Patrick and Christopher Harris (2000) "Strategic Experimentation: The Undiscounted Case," *Incentives, Organizations and Public Economics - Papers in Honour of Sir James Mirrlees ed. by P.J. Hammond and G. D. Myles*. Oxford: Oxford University Press, 53-68.

Bulow, Jeremy and Paul Klemperer (1994) "Rational Frenzies and Crashes," *Journal of Political Economy*, 102(1), 1-23.

Bulow, Jeremy and Paul Klemperer (1998) "Generalized War of Attrition," *American Economic Review*, 89(1), 175-189.

Folland, Gerald B. (1999) *Real Analysis: Modern Techniques and Their Applications*, Second Edition John Wiley & Sons, New York.

Fudenberg, Drew and Jean Tirole (1986) "A Theory of Exit in Duopoly," *Econometrica*, 54(4), 943-960.

Gallini, Nancy T. and Suzanne Scotchmer (2002) "Intellectual Property; When is it the Best Incentive Mechanism?" *Innovation Policy and the Economy*, 2, 51-78.

Gilbert, Richard. and Carl Shapiro (1990) "Optimal Patent Length and Breadth," *Journal of Industrial Economics*, 21(1), 106-112.

Gul, Faruk and Roger Lundholm (1995) "Endogenous Timing and Clustering of Agents' Decisions" *Journal of Political Economy*, 103(5), 1039-1066.

Haigh, John and Chris Cannings (1989) "The n -person War of Attrition," *Acta Applicandae Mathematicae*, 14(1-2), 59-74.

Hendricks, Kenneth, Andrew Weiss and Charles Wilson (1988) "The War of Attrition in Continuous Time with Complete Information," *International Economic Review*, 29(2), 663-680.

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

Kitch, E. W. (1977) "The Nature and Function of the Patent System," *The Journal of*

Law and Economics, 20(2), 265-290.

Klemperer, Paul (1990) “How Broad Should the Scope of Patent Protection Be?” *Journal of Industrial Economics*, 21, 113-130.

Maynard Smith, John (1974) “The Theory of Games and the Evolution of Animal Conflicts,” *Journal of Theoretical Biology*, 47(1), 29-19.

Reinganum, Jennifer (1981) “Dynamic Games of Innovation,” *Journal of Economic Theory*, 25, 21-24.

Scotchmer, Suzanne (1999) “On the Optimality of the Patent System,” *Rand Journal of Economics*, 30, 181-196.

Scotchmer, Suzanne (2005) *Innovations and Incentives*, MIT Press, Cambridge, Massachusetts.

Simon, Leo K. and Maxwell B. Stinchcombe (1989) “Extensive Form Games and Continues Time: Pure Strategies,” *Econometrica*, 57, 1171-1214.

Taylor, Curtis (1995) “Digging for the Golden Carrots: An Analysis of Research Tournaments,” *American Economic Review*, 85(4), 872-890.

Tirole, Jean (1988) *The Theory of Industrial Organization*, MIT Press, Cambridge, Massachusetts.