Learning by Selling, Knowledge Spillovers, and Patents*

Yves Guéron† — Seoul National University
Jihong Lee‡ — Seoul National University

January 2020

Abstract

We examine the incentives for experimentation in the context of innovation and market competition. A monopolist chooses whether to sell early-stage product or perform costly scale-up R&D. Early market participation facilitates learning about demand but invites knowledge spillovers and competitors, while R&D acts as a barrier to entry. We derive the firm's optimal policy and analyze the impact of various market characteristics. The model admits both under- and over-experimentation vis-à-vis the socially optimal policy. Patents can control the pace of innovation and restore the efficient level of experimentation. When the surplus from R&D is large, rewarding early-stage innovation limits socially wasteful investments.

Keywords: Innovation, entry deterrence, experimentation, knowledge spillovers, patent policy, petty patent

JEL codes: D83, L12, O31, O34

*The authors are grateful to Tilman Börgers, Qingmin Liu, Eeva Mauring, David Miller, Matt Mitchell, Caroline Thomas, and participants at a number of conferences and seminars for helpful comments. This research was supported by the Institute of Economic Research and the Park Yang Sook-Chung Yung Ho fund at Seoul National University.

†Department of Economics, Seoul National University, Seoul 08826, Korea; ygueron@snu.ac.kr

‡Department of Economics, Seoul National University, Seoul 08826, Korea; jihonglee@snu.ac.kr

Contents

1 Introduction

Dropbox, the cloud storage service that now boasts over 500 million users worldwide, began its public life meagerly as a short video showing consumers what it could do. After the video's release, the company's beta sign-up list jumped from 5,000 to 75,000 overnight. On the back of this level of interest, Dropbox proceeded with a full release of its product 6 months later. The rest is history.

The timeline of Dropbox illuminates a number of important characteristics of innovation and entrepreneurship that apply more generally. In particular, the five-minute video, released in March 2008 and already a classic case-study of the so-called "minimum viable product" strategy in business schools, enabled Dropbox to test if there was demand for its technology, and the rapid response gave it enough confidence to commit to further effort and financial resources.

Innovation and entrepreneurship are intimately connected via multiple layers of uncertainties. An innovating firm does not know whether its R&D investment will bear fruits or hit a dead-end (e.g. Arrow, 1962). Even after successful innovation, significant uncertainties remain with regard to the overall prospects and applications for the technology in the marketplace (e.g. Rosenberg, 1994). Indeed, how a firm navigates through the complex web of uncertainties holds key to its eventual prosperity.

Uncertainties naturally make experimentation a valuable component of entrepreneurship (e.g. Hayek, 1948; Rosenberg, 1994; Kerr, Nanda, and Rhodes-Kropf, 2014; Manso, 2016). While the literature on innovation and experimentation in economics has largely focused on uncertainties of *technological* nature, the case of Dropbox highlights the importance of *demand* uncertainty and learning by interacting with users (e.g. Lundvall, 1988; Von Hippel, 1988).

The IT revolution has drastically lowered the cost of market experimentation in software and many other industries as well as for start-up enterprises. But, as Kerr, Nanda, and Rhodes-Kropf (2014) point out, we still observe large differences across firms and industries in the degree of experimentation. For every Dropbox, there are numerous innovations that, despite hefty investments, falter as a result of lacking utility. Much-hyped Google Glass flopped not because of its technology but rather its failure to 'offer anything that average people really want, let alone need, in their everyday lives'.[1]

[1] See https://slate.com/technology/2014/11/google-glass-why-its-time-for-google-to-walk-away-from-its-much-hyped-moonshot.html

How can we explain the observed heterogeneity in market experimentation and the contrasting fortunes of Dropbox and Google Glass? What are the trade-offs associated with market experimentation and their welfare implications? Is there a role for policy intervention, and if so, how?

To address these questions, we set up a simple theoretical model of market competition in the framework of experimentation with exponential bandits. A monopolist possesses a piece of novel technology but faces uncertain demand. It chooses to enter the market early with a product at the beginning of its development, or continue investing in R&D until the technology matures. Early market participation, or "learning by selling", offers informational benefits but invites imitation and entry from competitors. By delaying commercialization and investing in further "scale-up" R&D, the monopolist improves the likelihood of a "quiet life" (Hicks, 1935).[2]

Specifically, a firm in continuous time chooses whether to experiment or engage in costly R&D. If demand is bad, no sales can occur. If demand is good and the firm experiments, sales arrive randomly. While experimentation facilitates learning, it attracts entrants. Upon entry, competition drives profits to zero and the incumbent exits the market. The success of R&D, which also arrives randomly, not only improves the value of the innovation but acts as a barrier to entry.

The optimal policy of the firm exhibits threshold dynamics, choosing to invest in R&D when sufficiently optimistic about demand and to experiment otherwise. While experimenting, the first sale fully reveals the good state of demand and the firm finds it optimal to invest in R&D; if no sales arrive, the firm becomes increasingly pessimistic. Experimentation takes place over a wider range of beliefs as the surplus from R&D falls. Faster learning or delayed arrival of successful R&D also has a positive effect on experimentation. But, as knowledge diffuses more rapidly, it becomes more important to pursue R&D and reach a higher technological stage.

We then analyze the social incentives for experimentation by considering the optimal policy of a planner whose objective incorporates the positive externalities from knowledge spillovers. Two types of laissez-faire inefficiency are identified. *Under*-experimentation occurs when the surplus from R&D is high, and the firm proceeds with innovation too rapidly in the sense that it performs R&D at beliefs where the social planner would experiment in the market

[2] Hicks (1935, p.8) famously remarked that 'the best of all monopoly profits is a quiet life'.

instead. *Over*-experimentation occurs, and the firm accumulates innovation too slowly, when the surplus from R&D is low.

The private and social benefits of experimentation diverge in the above manner because of the multiple effects of knowledge spillovers. On the one hand, R&D is costly, but the society benefits whether it is undertaken by the incumbent or an entrant. Thus, one should expect under-investment in R&D, or over-experimentation. But, on the other hand, R&D helps to protect the monopolist's profits from imitators. This entry deterrence effect of R&D weakens the value of experimentation to the firm vis-à-vis the society. The overall effect depends on the surplus from R&D.

Finally, we turn to the issue of policy intervention by asking the role of intellectual property rights. This normative question is of particular relevance as there are even suggestions of the government's own hand in under-utilization of information and excessive commitment of public resources to R&D (e.g. Nelson, 1977; Rosenberg, 1994).

A new trade-off is introduced. At each stage of the technology, the monopolist can protect its innovation with a patent, or resort to trade secret. A patent shuts down entry for a given length of time but intensifies the spillover effects such that, upon expiry, imitators arrive at a faster rate. As noted by Rosenberg (1994), even if the patent does not disclose sufficient details to trigger imitation, the mere knowledge of feasibility may be valuable itself.[3]

We find a patent policy that aligns private with social incentives for experimentation. The proposed policy is based on a two-tier patent system and depends on the surplus from R&D. Specifically, for an over-experimenting monopolist with low-surplus innovation, regular patent offered at the mature stage of technology restores socially optimal behavior, and under-experimentation, which occurs with high-surplus innovation, can be corrected by introducing patent protection at the early stage of technology. One can interpret the latter type of incentives as "petty" patents with reduced "patentability" standards.[4]

Our analysis thus suggests a new role for patents, whose long existence has largely been associated with incentive provision for a public good. In our model, patents directly serve to control the *pace* of innovation. In particular, by blocking competition from imitators, petty patent enables the high-surplus firm to experiment in the market at beliefs where it would

[3]The spillover effects of patents, as well as the benefits of trade secret, have been scrutinized by a number of studies. See, for example, Cohen, Nelson, and Walsh (2000) and Arundel (2001).

[4]Petty patents are available in many countries under various definitions, though notably not in US, UK, or Canada. See Section 5.4 for more discussion.

otherwise hasten the pursuit of next stage innovation, and hence, save on socially wasteful R&D expenditure.

Our two-step innovation setup resembles the sequential innovation models of, among many others, Green and Scotchmer (1995), Denicoló (2000), and Hopenhayn, Llobet, and Mitchell (2006). In this body of research, the prospect of "creative destruction" generates an additional hold-up problem for the first innovator, who should therefore be compensated to take on the initial investment. In our model, the threat of competitive entry affects the incumbent's learning decision, and rewarding early-stage innovation serves instead to stimulate experimentation and slow down the pace of innovation. Our results contrast also with O'Donoghue (1998), who argues for a stronger patentability requirement to promote R&D by delaying the follow-on patentable innovation.

While the existing studies are mainly concerned with linkages across separate innovations, our model can be interpreted as describing incremental steps within a single innovation. To highlight the trade-off between investment and experimentation, we rule out Schumpeterian forces. The intrinsic nature of R&D in our model remains identical irrespective of the performing entity. This makes our policy objective different from the sequential innovation literature, whose focus is on achieving optimal policy mix to balance protection of early innovations to stimulate investment with deadweight losses from delayed arrival of superior technologies.

One can view our paper as introducing experimentation to the long-standing literature on market competition, and especially, strategic entry deterrence via supply-side commitments, which dates back to Spence (1977). As in Gilbert and Newbery (1982), the incumbent in our model wants to preempt entrants with R&D investment, but only in the good state of demand. Learning invites knowledge spillovers and competitive entry, in the spirit of technology leakage first suggested by Arrow (1962). These two classic effects, by Spence and Arrow, are brought together by market uncertainty and lie at the core of the under- and over-experimentation phenomena, yielding novel policy implications.[5]

There is also an abundance of so-called "innovation race" models, initiated by Loury (1979) and Dasgupta and Stiglitz (1980). While learning was introduced to this early literature by Choi (1991),[6] the more recent exponential bandit model of experimentation, due to Keller,

[5]Among the early studies on market competition, Reinganum (1983) also finds that the extent of inefficient capital accumulation is linked to the surplus from R&D. In a recent paper, Parra (2019) revisits this literature in a sequential innovation setup and explores the role of patent policy, but there is no learning.

[6]See also Malueg and Tsutsui (1997) and Weeds (2002).

Rady, and Cripps (2005), offers a richer framework to study dynamic R&D competition and related issues, and has been extended to many other strategic settings (e.g. Keller and Rady, 2010, 2015; Strulovici, 2010; Bonatti and Hörner, 2011; Klein and Rady, 2011; Murto and Välimäki, 2011).[7]

Among these, our paper is most closely related to Keller and Rady (2015). In this model, each player chooses between risky and safe arms. The safe arm generates a known constant flow cost while the risky arm could result in a large lump-sum cost, or "breakdown", arriving randomly only in one state of nature. In equilibrium, sufficiently optimistic players dangle with the risky arm until the breakdown occurs, in which case they switch to the safe arm. No learning takes place below the belief threshold.

The equilibrium dynamics in our single-player model are similar,[8] but our setup differs in two important aspects. First, the learning arm in our model (i.e. selling) receives multiple news (i.e. entry and sales), triggering exit as well as switch to the non-learning arm (i.e. R&D). Second, the value of R&D is itself uncertain.[9] In terms of results, the focus of Keller and Rady (2015) and the corresponding literature is the externality that an individual's learning imposes on others.[10] This causes under-experimentation. The inefficiency in our model is driven instead by the interaction between two mechanisms: the discrepancy between private and social returns to knowledge spillovers on the one hand, and the strategic use of R&D as entry-deterring device on the other. As a consequence, we observe both under- and over-experimentation.

Policy design issues to address the under-provision problem are taken up by Akcigit and Liu (2015) and Halac, Kartik, and Liu (2017) in innovation race settings with exponential bandits. Unlike our paper, they consider one-off innovation with technological uncertainty as well as policy instruments other than patents.[11] There is also a growing literature on contracting with single experimenting agent and moral hazard (e.g. Bergemann and Hege, 1998, 2005; Manso, 2011; Hörner and Samuelson, 2013; Halac, Kartik, and Liu, 2016; Bonatti and Hörner, 2017). The source of inefficiency in these models is the payoff discrepancy in a

[7] Bolton and Harris (1999) propose an alternative framework based on Brownian motion.

[8] As in Keller and Rady (2015), the value function in our model consists of two linear segments and does not involve smooth pasting.

[9] Camargo (2007) examines correlated arms in another bandit framework.

[10] Besanko and Wu (2013) add direct externality effects from knowledge spillovers in one of the papers.

[11] Acemoglu, Bimpikis, and Ozdaglar (2011) consider patents to spur R&D spending in an innovation race setting with simple learning structure.

principal-agent relationship, but in contrast to our setup, the agent's bias is fixed and the principal faces only under-experimentation.[12]

The rest of the paper is organized as follows. Section 2 describes the model of market competition and experimentation by a monopolist. In Section 3, we characterize the optimal policy of the monopolist. The divergence between private and social incentives for experimentation are then identified and explained in Section 4. Section 5 discusses the role of patents as a tool for correcting inefficient experimentation, including specific policy implications. Section 6 concludes by summarizing our results and connecting them to the motivating examples. Appendix contains some of the analyses and proofs left out from the main text for expositional reasons.

2 Model

We consider a monopolist in possession of a novel technology and operating in continuous time, $t \in [0, \infty)$, with real discount rate $r > 0$. The firm can commercialize the technology and sell the product across two "stages", $s \in \{0, 1\}$, of its development. The cost of production in either stage is normalized to zero.

The market demand can be in one of two "states", $\theta \in \{0, 1\}$. When $\theta = 0$, or the state is "bad", there is no demand for the product whether $s = 0$ or $s = 1$. The firm does not know the demand state, and the prior belief on $\theta = 1$, or "good" state with positive demand, is denoted by $\mathbb{P}(\theta = 1) = p_0$. Let $p(t)$ denote the corresponding belief at time t.

Initially, $s = 0$, or the technology is in its "early" stage. At every t, while remaining in $s = 0$, the firm chooses to either experiment by commercializing an early version of the technology or further invest in R&D to scale up the technology to the "mature" stage $s = 1$. When experimenting in the market, or "learning by selling", the firm receives no sales if demand is bad; if demand is good then sales arrive randomly according to an exponential distribution with parameter $\lambda_0 > 0$.

If the firm chooses R&D at any t, it pays an immediate sunk cost $c > 0$ and experimentation stops. R&D is completed after a random time, which also follows an exponential distribution with parameter $\kappa > 0$. Upon completion, the technological stage advances from $s = 0$ to

[12] Guo (2016) obtains both types of suboptimal experimentation in a model of optimal delegation with adverse selection. In this model, the agent's bias is also fixed, but under- or over-experimentation occurs depending on the realization of hidden information.

$s = 1$ and the firm sells its product. The arrival rate of sales in $s = 1$, conditional on $\theta = 1$, is given by $\lambda_1 > 0$, which may differ from λ_0.[13]

Market participation generates knowledge spillovers: competitors can potentially copy the innovation and enter the market. Specifically, in each $s \in \{0, 1\}$, a continuum of firms enter, and the incumbent exits, after a random time distributed exponentially with parameter $\mu_s > 0$, where $\mu_0 > \mu_1$.[14] The entry probability itself depends on the stage of technological development, and moreover, R&D acts as a barrier to entry. We refer to the latter feature as the "entry deterrence effect of R&D".

Let w_s denote the total surplus generated by the innovation in stage $s \in \{0, 1\}$, where $w_1 > w_0 = 0$. The R&D investment raises the value of the innovation, and the total surplus in $s = 0$ is normalized to zero. We also assume that the firm can extract all the surplus when operating as monopolist. This simplifies the analysis and, as will be clarified below, helps to crystalize the effects of knowledge spillovers in our model.[15]

Remark 1. While the exogenous parameters w_s and λ_s correspond to information on market demand, the supply characteristics are captured by c and κ. The cost of R&D, c, represents the relative price of experimentation emphasized in IT industries; κ is another component of R&D which varies across markets in meaningful ways.[16] The parameter μ_s measures the speed of knowledge spillovers, which is affected by frictions such as the distance between inventors (e.g. Jaffe, Trajtenberg, and Henderson, 1993). Regarding the assumption that $\mu_0 > \mu_1$, Aghion et al. (2009) report evidence of causal relationship between entry threat and incumbent innovation.

Belief Updating Consider any technological stage $s \in \{0, 1\}$, belief $p(t)$ at time t, and time interval $[t, t + dt)$. If sales arrive during the interval then the belief jumps to $p(t + dt) = 1$. Suppose now that no sales occur. If $\theta = 0$, this happens with probability 1; if $\theta = 1$, the corresponding probability is $e^{-\lambda_s dt} \simeq 1 - \lambda_s dt$. Thus, conditional on no sales occurring, beliefs

[13]Since R&D is irreversible, allowing for selling and learning during the R&D phase will not affect the results. In $s = 1$, learning does take place but there is no alternative decision to make.

[14]In $s = 0$, competition means that the costly R&D can be taken up by only one firm. With a continuum of firms, the probability of the incumbent remaining in the market is zero even if demand is good.

[15]See Remark 3.

[16]The relevance of κ is advocated by Budish, Roin, and Williams (2015), who study patent policy in the pharmaceutical industry where the length of clinical trials exhibits large variations.

are updated as follows according to Bayes' rule:

$$p(t+dt) = \frac{(1-\lambda_s dt)p(t)}{(1-\lambda_s dt)p(t) + (1-p(t))} = \frac{(1-\lambda_s dt)p(t)}{1-\lambda_s dt p(t)}.$$

Subtracting $p(t)$ from both sides, dividing by dt, and letting $dt \to 0$, we find the following law of motion for beliefs, conditional on no sales arriving, which is standard in exponential bandit models:

$$\frac{dp(t)}{dt} = -\lambda_s p(t)\left[1-p(t)\right]. \tag{1}$$

3 Incentives for Experimentation

3.1 Dynamic Programming

We begin by setting up the dynamic programming problem of the monopolist. Let $U_s(p)$ denote the value function of the firm when the technological stage is $s \in \{0,1\}$ and the belief on the good demand state is p.

In $s=1$, there is no further action to take for the monopolist, and it is straightforward to pin down the corresponding continuation value, which is linear in p. For $p=1$, the continuation value simply amounts to the expected profits from the random flows of sales, discounted by the probability of competitive entry.

Lemma 1. *We have*

$$U_1(p) = pU_1(1) = \frac{p\lambda_1}{\mu_1 + r} w_1.$$

Proof. First, consider a small interval of time dt in $s=1$ and $p=1$:

- A sales arrive with probability $1-e^{-\lambda_1 dt}$, yielding the instantaneous payoff w_1;
- The continuation value $U_1(1)$ is obtained if there is no entry, which occurs with probability $1-e^{-\mu_1 dt}$.

Thus, we have

$$U_1(1) = (1-e^{-\lambda_1 dt})w_1 + e^{-rdt}e^{-\mu_1 dt}U_1(1).$$

Using the the first-order approximations $1-e^{-\lambda_1 dt} \simeq \lambda_1 dt$ and $e^{-rdt}e^{-\mu_1 dt} \simeq 1-(r+\mu_1)dt$, and letting $dt \to 0$, it follows that

$$U_1(1) = \frac{\lambda_1}{\mu_1 + r} w_1.$$

Second, note that we must have $U_1(p) = pU_1(1)$. This is because no further action can be taken by the firm in $s = 1$. With probability p, demand is good, in which case the value function is $U_1(1)$, and with the complementary probability, demand is bad, in which case the value is zero. This completes the proof. □

Next, consider $s = 0$ and belief p. If the firm decides to pursue R&D, it must sink the immediate cost c and wait on average $1/\kappa$ units of time for R&D to be completed. There is no learning during this interval. The corresponding payoff is then given by

$$-c + \frac{\kappa}{\kappa + r} U_1(p).$$

Suppose that the firm experiments by selling in the market, and consider a small interval of time dt. Three events can occur:

- Sales arrive with probability $p(1 - e^{-\lambda_0 dt})$, in which case the belief jumps to $p = 1$ and the firm engages in R&D; therefore, the value is

$$-c + \frac{\kappa}{\kappa + r} U_1(1);$$

- Entry arrives with probability $(1 - e^{-\mu_0 dt})$, in which case the value is zero;
- Neither sales nor entry arrives with the complementary probability $\left[1 - p(1 - e^{-\lambda_0 dt})\right] e^{-\mu_0 dt}$, in which case the belief drifts down to $p + dp$ and the continuation value is $U_0(p + dp)$, discounted by e^{-rdt}.

We therefore obtain the following Bellman equation for $U_0(p)$.

Lemma 2. *We have*

$$U_0(p) = \max \left\{ -c + \frac{\kappa}{\kappa + r} pU_1(1), \lambda_0 p dt \left[-c + \frac{\kappa}{\kappa + r} U_1(1) \right] \right.$$
$$\left. + \left[1 - (\lambda_0 p + \mu_0 + r)dt\right] U_0(p + dp) \right\}, \quad (2)$$

where, by Lemma 1,

$$U_1(1) = \frac{\lambda_1}{\mu_1 + r} w_1.$$

This function is convex and continuous on $p \in [0, 1]$.

Proof. Given the arguments preceding the claim, it suffices to prove the convexity and continuity of (2).

Consider the following two-stage procedure. First, nature decides on the probability with which demand is good. With probability β, demand is good with probability p, and with the complementary probability the good state occurs with probability q. In the second stage, nature decides whether demand is good or bad. The value function for this problem is $U_0(\beta p + (1-\beta)q)$.

Suppose now that the firm is informed about the outcome of the first stage. The value function in that case is $\beta U_0(p) + (1-\beta)U_0(q)$. Since it must be better to be informed than uninformed, we must have

$$\beta U_0(p) + (1-\beta)U_0(q) \geq U_0\left[\beta p + (1-\beta)q\right],$$

implying convexity of $U_0(p)$. □

3.2 Optimal Policy

Suppose that demand is known to be good (i.e. $p = 1$) but the monopolist is yet to invest in R&D. Then, by Lemma 1, the continuation value of R&D, and advancing to $s = 1$, is

$$-c + \frac{\kappa}{\kappa + r} U_1(1) = -c + \frac{\kappa}{\kappa + r}\frac{\lambda_1}{\mu_1 + r} w_1.$$

The continuation value from remaining in $s = 0$ indefinitely is normalized to 0. Thus, when $p = 1$, the monopolist invests in R&D if and only if

$$c \leq \bar{c} := \frac{\kappa}{\kappa + r}\frac{\lambda_1}{\mu_1 + r} w_1. \tag{3}$$

The firm's optimal policy is in cutoff with the corresponding value function exhibiting piecewise linearity and no smooth pasting, similarly to Keller and Rady (2015).

Proposition 1. *For $c < \bar{c}$, where $\bar{c}$ is given by (3), the optimal policy of the monopolist is to experiment below the cutoff belief*

$$\hat{p} = \frac{(\lambda_0 + \mu_0 + r)c}{\frac{\kappa}{\kappa + r}(\mu_0 + r)U_1(1) + \lambda_0 c} \in (0, 1)$$

and invest in R&D above. For $c \geq \bar{c}$, $\hat{p} = 1$.

The value function, $U_0(p)$, is continuous, strictly increasing, and piecewise linear with a single kink at $\hat{p}$. We have

$$U_0(p) = \begin{cases} -c + \frac{\kappa}{\kappa+r} p U_1(1) & \textit{if} \quad p \geq \hat{p} \\ \frac{\lambda_0 p}{\lambda_0+\mu_0+r}\left[-c + \frac{\kappa}{\kappa+r} U_1(1)\right] & \textit{otherwise,} \end{cases}$$

where, by Lemma 1,

$$U_1(1) = \frac{\lambda_1}{\mu_1 + r} w_1.$$

Proof. Note first that $c < \bar{c}$ ensures that $\hat{p} \in (0, 1)$. If $c \geq \bar{c}$, it is not desirable to invest in R&D, irrespective of the belief, and hence, $\hat{p} = 1$.

We show that the function $U_0(p)$ stated in the claim solves the Bellman equation (2) in Lemma 2 and is piecewise linear since $U_1(1)$ is constant (Lemma 1). It then corresponds to the value function of the firm, and the proposed policy achieves the maximum in (2) for any belief p.

When it is optimal to experiment instead of R&D, (2) becomes

$$U_0(p) = \lambda_0 p dt \left[-c + \frac{\kappa}{\kappa + r} U_1(1)\right] + [1 - (\lambda_0 p + \mu_0 + r)dt]\, U_0(p + dp).$$

Subtracting $U_0(p + dp)$ from both sides, dividing by $dt = -dp/[\lambda_0 p(1 - p)]$ as in (1), and letting $dp \to 0$, we obtain the following first-order ordinary differential equation:

$$\lambda_0 p(1 - p)U_0'(p) + (\lambda_0 p + \mu_0 + r)U_0(p) = \lambda_0 p \left[-c + \frac{\kappa}{\kappa + r} U_1(1)\right]. \tag{4}$$

Given the initial condition $U_0(0) = 0$, It is straightforward to show that (4) has the following solution:

$$\frac{\lambda_0 p}{\lambda_0 + \mu_0 + r}\left[-c + \frac{\kappa}{\kappa + r} U_1(1)\right].$$

The cutoff $\hat{p}(c)$ is found by equating the solution to (4) with the continuation value from R&D, i.e.

$$-c + \frac{\kappa}{\kappa + r} p U_1(1),$$

and then substituting for $U_1(1) = \frac{\lambda_1}{\mu_1+r} w_1$ as in Lemma 1. □

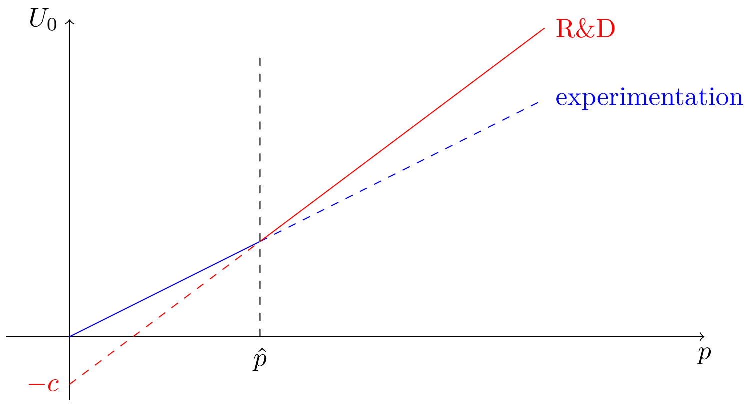

Figure 1: Value function

To better understand the monopolist's incentives (and later compare them to those of a social planner), we next discuss in detail the structure of the value function, which is illustrated in Figure 1. Let us re-write the value function as

$$U_0(p) = \begin{cases} U_0^{rnd}(p) = -c + \hat{k}_1 p & \text{if} \quad p \geq \hat{p}(c) \\ U_0^{exp}(p) = \left(\hat{k}_2 - \hat{k}_3 c\right) p & \text{otherwise,} \end{cases}$$

where

$$\hat{k}_1 = \frac{\kappa}{\kappa + r} \frac{\lambda_1}{\mu_1 + r} w_1; \tag{5}$$

$$\hat{k}_2 = \frac{\lambda_0}{\lambda_0 + \mu_0 + r} \left(\frac{\kappa}{\kappa + r} \frac{\lambda_1}{\mu_1 + r} w_1 \right); \tag{6}$$

$$\hat{k}_3 = \frac{\lambda_0}{\lambda_0 + \mu_0 + r}. \tag{7}$$

The coefficient $\hat{k}_1$ represents the marginal benefit from an increase in p when performing R&D. It consists of, conditional on demand being good, the expected profits in the mature technological stage $s = 1$, discounted by the length of completing R&D and the possibility of competitive entry thereafter.

The marginal effect from an increase in p when experimenting consists of two terms. The coefficient $\hat{k}_2$ represents the delayed benefit from R&D, which is also discounted by the possibility of entry in $s = 0$. The coefficient $\hat{k}_3$ on the cross term $c \cdot p$ in $U_0^{exp}(p)$ captures the cost. A higher p means a greater probability with which sales arrive, and this in turn raises the likelihood of incurring the fixed cost c of R&D.

Given this re-parametrization, the monopolist's belief threshold can be re-written as

$$\hat{p} = \frac{c}{\hat{k}_1 - \hat{k}_2 + \hat{k}_3 c}.$$

We then see that an increase in $\hat{k}_1 - \hat{k}_2$ lowers the threshold $\hat{p}$ and expands the region of beliefs in which the firm invests in R&D. This is intuitive, as $\hat{k}_1 - \hat{k}_2$ measures the marginal benefit that greater confidence in market demand has on R&D relative to experimentation. An increase in $\hat{k}_3$ also lowers $\hat{p}$.

As noted, for instance, by Kerr, Nanda, and Rhodes-Kropf (2014), we observe substantial differences in the degree of experimentation across markets. Let us summarize how the incentives for experimentation depend on market characteristics by offering some comparative static results on $\hat{p}$.

Corollary 1. *For $c < \bar{c}$, we obtain the following:*

(i) $\frac{\partial \hat{p}}{\partial w_1} < 0$, $\frac{\partial \hat{p}}{\partial c} > 0$, *and* $\frac{\partial \hat{p}}{\partial \kappa} < 0$;

(ii) $\frac{\partial \hat{p}}{\partial \lambda_0} > 0$ *and* $\frac{\partial \hat{p}}{\partial \lambda_1} < 0$;

(iii) $\frac{\partial \hat{p}}{\partial \mu_0} < 0$ *and* $\frac{\partial \hat{p}}{\partial \mu_1} > 0$.

Proof. See Appendix A. □

Parts (i) and (ii) imply that the firm's incentives for experimentation are determined by the expected surplus from R&D (i.e. $\mathbb{E}[w_1 - c]$). As for the effects of its component parts, first, as w_1 falls or c rises, the surplus from R&D decreases and the optimal cutoff belief increases. Similarly, the firm engages in market experimentation over a wider range of beliefs also when it takes longer to complete the R&D (i.e. $\downarrow \kappa$), or when sales arrive in the mature stage more slowly relative to the early stage of technology (i.e. $\uparrow \lambda_0$ or $\downarrow \lambda_1$).

Another important market characteristics are the rates of competitive entry, μ_0 and μ_1. In part (iii), we first see that $\hat{p}$ decreases with μ_0. As knowledge diffuses more rapidly in $s = 0$, it becomes more important for the firm to scale up and reach $s = 1$. In contrast, $\hat{p}$ increases with μ_1. The difference between μ_0 and μ_1 measures the entry deterrence effect of R&D. Note that this effect works in the same direction as the surplus effect described above.

The incentives for experimentation weaken as the entry-deterrence effect of R&D strengthens, as is the effect of greater surplus from R&D.

4 Inefficient Pace of Innovation

4.1 Social Incentives for Experimentation

Overview Innovation is a complex phenomenon with multiple trade-offs. In this paper, we focus on the potential disparity between private and social incentives for experimentation that arises due to knowledge spillovers.

We consider a benevolent social planner who solves the problem of optimal public provision of the product in our dynamic setting with demand uncertainty and potential entry. While entry erodes the monopolist's profits, the society continues to benefit. We assume that this happens via two channels. First, in both $s = 0$ and $s = 1$, the social planner's post-entry continuation value is positive, involving the random flow of the surplus w_s. Second, when entry occurs in $s = 0$, uncertainty is immediately revealed, and if $\theta = 1$, one of the firms receives a sale and engages in R&D.

While the first channel echoes the standard externality problem associated with R&D, the second channel, the learning effect of entry, is new. As more firms enter the market and experiment with early-stage products, the likelihood of successful learning increases; with a continuum of such firms, the probability that at least one firm receives a sale, conditional on good demand, is 1 at any given time.

The informational benefits from knowledge spillovers accrue only in $s = 0$; in $s = 1$, entry does not affect the conditional distribution of sales. While this assumption simplifies the analysis and our results below are robust to alternative treatments of the consequences of entry (see Remark 2), the present formulation adopts the following perspective on the nature of supply and demand. When the technology is new, more firms and greater resources spent on experimentation result in faster arrival of consumers. Supply creates demand *à la* Say's law. When the technology is mature, however, demand is less dependent on the market structure since consumers are likely to possess greater awareness of the product by this stage.

Optimal Policy Let $W_s(p)$ denote the social value function given technological stage s and belief p. This can be readily derived from the analysis of Section 3 by making the following two modifications. First, since the society is not concerned by entry, we can delete both μ_0 and μ_1 in the formulation of dynamic programming for the social planner. Second, we assume that entry leads to immediate arrival of a sale if $\theta = 1$ and can therefore set the corresponding arrival rate in the planner's problem as $\lambda_0 + \mu_0$, instead of λ_0 as in the firm's problem.

When $p = 1$, the social value of R&D amounts to

$$-c + \frac{\kappa}{(\kappa + r)} \frac{\lambda_1}{r} w_1,$$

implying that the planner prefers to invest in R&D and move onto $s = 1$ rather than indefinitely remain in $s = 0$ if and only if

$$c \leq \bar{c}^* := \frac{\kappa}{(\kappa + r)} \frac{\lambda_1}{r} w_1. \tag{8}$$

But, notice that $\bar{c}^* > \bar{c}$, where $\bar{c}$ is given by (3). The upper bound on the range of costs in which the society can potentially benefit from R&D is greater than that for the firm. This reflects the fact that the society continues to benefit even after the arrival of entrant in $s = 1$. Due to this externality, the planner would be willing to incur a greater cost than the monopolist to scale up.

We now state the optimal policy of the planner, which can be derived directly from Proposition 1 with the two modifications mentioned above.

Proposition 2. *For $c < \bar{c}^*$, where $\bar{c}^*$ is given by (8), the optimal policy of the social planner is to experiment below the cutoff belief*

$$p^* = \frac{(\lambda_0 + \mu_0 + r)c}{\frac{\kappa}{\kappa + r} r W_1(1) + (\lambda_0 + \mu_0)c} \in (0, 1)$$

and invest in R&D above. For $c \geq \bar{c}^*$, $p^* = 1$.

The value function, $W_0(p)$, is continuous, strictly increasing, and piecewise linear with a single kink at p^*. We have

$$W_0(p) = \begin{cases} -c + \frac{\kappa}{\kappa + r} p W_1(1) & \text{if} \quad p \geq p^* \\ \frac{(\lambda_0 + \mu_0)p}{\lambda_0 + \mu_0 + r} \left[-c + \frac{\kappa}{\kappa + r} W_1(1) \right] & \text{otherwise,} \end{cases}$$

where $W_1(1) = \lambda_1 w_1 / r$.

Comparative statics on the social belief threshold, p^*, are reported below.

Corollary 2. *For $c < \bar{c}^*$, we obtain the following:*

(i) $\frac{\partial p^*}{\partial w_1} < 0$, $\frac{\partial p^*}{\partial c} > 0$, *and* $\frac{\partial p^*}{\partial \kappa} < 0$;

(ii) $\frac{\partial p^*}{\partial \lambda_0} > 0$ *and* $\frac{\partial p^*}{\partial \lambda_1} < 0$;

(iii) $\frac{\partial p^*}{\partial \mu_0} > 0$ *and* $\frac{\partial p^*}{\partial \mu_1} = 0$.

Proof. See Appendix A. □

These results further highlight the difference between private and social incentives for experimentation. They differ from Corollary 1 in part (iii). Due to the learning effect, entry in the early stage of technology is beneficial for welfare. Thus, the social threshold, p^*, increases, and the planner experiments over a wider range of beliefs, as entrants arrive at a faster rate in $s = 0$. Moreover, the entry rate in $s = 1$ is irrelevant for the planner. This implies that the entry-deterrence effect of R&D plays no role in determining the socially optimal level of experimentation.

4.2 Under- and Over-experimentation

To compare the optimal policies of the firm and the planner, we introduce the notions of under- and over-experimentation. The case of under-experimentation is illustrated in Figure 2.

Definition 1. We say that there is *under-experimentation* if $\hat{p} < p^*$ and $p \in (\hat{p}, p^*)$ and *over-experimentation* if $\hat{p} > p^*$ and $p \in (p^*, \hat{p})$.

One can interpret under-experimentation as the firm pursuing sequential innovation at a faster *pace* than the social planner as it undertakes scale-up R&D over a region of beliefs where the planner would instead experiment in the market. Similarly, over-experimentation can be thought of as slower than efficient pace of technological progress, paralleling the classic case of under-investment in R&D due to positive externalities.

Our model admits both under- and over-experimentation. As noted above, there are two main determinants of the incentives for experimentation in our model: the surplus from R&D and the entry-deterrence effect of R&D. While the second effect is relevant only to the firm, the extent of the first effect differs across the firm and the planner due to the externalities from knowledge spillovers. The direction of inefficiency depends on the size of the surplus from R&D, which we choose to measure by its cost c throughout below.

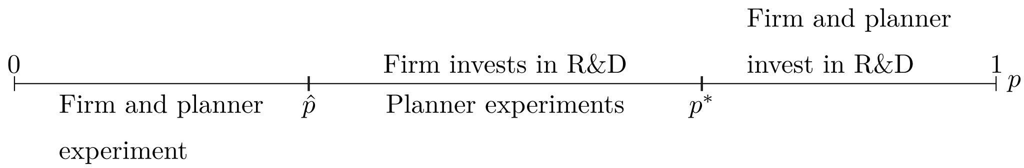

Figure 2: Under-experimentation

Proposition 3. *There exists* $\gamma \in (0, \bar{c})$ *such that* $\hat{p}(c) < p^*(c)$ *for any* $c < \gamma$ *and* $\hat{p}(c) > p^*(c)$ *for any* $c > \gamma$*. The threshold* γ *is given by*

$$\gamma = \frac{\kappa}{\kappa + r} \frac{\lambda_1(\mu_0 - \mu_1)}{\mu_0(\mu_1 + r)} w_1.$$

Proof. First, it is straightforward to derive from the expressions of $\hat{p}(c)$ and $p^*(c)$ given in Propositions 1 and 2 that $\hat{p}(0) = p^*(0) = 0$ and $\hat{p}'(0) < p^{*\prime}(0)$. Thus, for c sufficiently close to zero, $\hat{p}(c) < p^*(c)$.

Next, note that $\hat{p}(c) = 1$ if and only if $c \geq \bar{c}$ and $p^*(c) = 1$ if and only if $c \geq \bar{c}^*$. Since $\bar{c}^* > \bar{c}$, it follows that $\hat{p}(c) > p^*(c)$ for c sufficiently close to $\bar{c}$ and the two thresholds must cross at some $c < \bar{c}$. The intersection γ is found by solving $\hat{p}(\gamma) = p^*(\gamma)$. □

Figure 3 plots the thresholds $\hat{p}(c)$ and $p^*(c)$, which are both increasing but with different slopes. The private and social thresholds cross at $c = \gamma < \bar{c}$.[17] Thus, we have both under- and over-experimentation phases. When the cost of R&D is low, the threshold of the firm is below the threshold of the planner and the region between the two thresholds exhibits under-experimentation. When c is close to $\bar{c}$, we have $\hat{p} > p^*$, and at $p \in (p^*, \hat{p})$, the firm over-experiments. For any $c \in (\bar{c}, \bar{c}^*)$, the monopolist never finds it optimal to invest in R&D, while the planner does so with sufficiently optimistic beliefs.

We summarize below how the intersection between $\hat{p}(c)$ and $p^*(c)$ is determined by the main parameters of the model.

Corollary 3. *We have the following:*

(i) $\frac{\partial \gamma}{\partial w_1} > 0$ *and* $\frac{\partial \gamma}{\partial \kappa} > 0$;

(ii) $\frac{\partial \gamma}{\partial \lambda_0} = 0$ *and* $\frac{\partial \gamma}{\partial \lambda_1} > 0$;

[17] To save space, Figure 3 truncates the region containing $\bar{c}^*$.

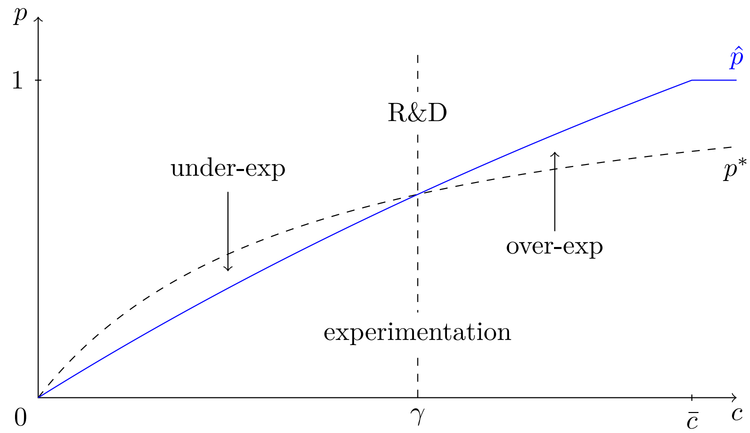

Figure 3: Social (p^*) and private ($\hat{p}$) thresholds

(iii) $\frac{\partial \gamma}{\partial \mu_0} > 0$ *and* $\frac{\partial \gamma}{\partial \mu_1} < 0$.

Proof. This is straightforward. □

These comparative statics on γ reveal that the under-experimentation phenomenon becomes more prevalent as R&D becomes more valuable (as $\uparrow w_1$, $\uparrow \kappa$, or $\uparrow \lambda_1$), or as the entry-deterrence effect of R&D strengthens (as $\uparrow \mu_0$ or $\downarrow \mu_1$).

Sources of Inefficiency We now discuss in detail the sources of inefficient experimentation in our model. In particular, we pay close attention to how both under- and over-experimentation emerge due to the differential impact that knowledge spillovers have on private and social returns to R&D, respectively.

Let us re-write the value function of the planner as

$$W_0(p) = \begin{cases} W_0^{rnd}(p) = -c + k_1^* p & \text{if} \quad p \geq \hat{p}(c) \\ W_0^{exp}(p) = (k_2^* - k_3^* c)p & \text{otherwise,} \end{cases}$$

where

$$k_1^* = \frac{\kappa}{\kappa + r} \frac{\lambda_1 w_1}{r};$$
$$k_2^* = \frac{\lambda_0 + \mu_0}{\lambda_0 + \mu_0 + r} \left(\frac{\kappa}{\kappa + r} \frac{\lambda_1 w_1}{r} \right);$$
$$k_3^* = \frac{\lambda_0 + \mu_0}{\lambda_0 + \mu_0 + r}.$$

These coefficients can be interpreted similarly to the corresponding coefficients, $\hat{k}_1$, $\hat{k}_1$, and $\hat{k}_3$, in the monopolist's value function as in (5)-(7), and the planner's belief threshold can be re-written as

$$p^*(c) = \frac{c}{k_1^* - k_2^* + k_3^* c}.$$

Comparing p^* with $\hat{p}$, it is then straightforward to verify that

(i) $\hat{p}(0) = p^*(0) = 0$;

(ii) $\hat{p}'(0) = \frac{1}{\hat{k}_1 - \hat{k}_2} < \frac{1}{k_1^* - k_2^*} = p^{*\prime}(0)$.

This implies that, for low values of c, the private threshold $\hat{p}(c)$ must lie below the social threshold $p^*(c)$ and our model must exhibit *under*-experimentation, as demonstrated Figure 3. When it is costless, R&D improves the value of the innovation for both the firm and the planner whenever $p > 0$ (part (i)), but the marginal benefit is greater for the firm (part (ii)).

The reason for the latter part can be seen from

$$\begin{aligned} & \hat{k}_1 - \hat{k}_2 > k_1^* - k_2^* \\ \Leftrightarrow \quad & \frac{1}{\lambda_0 + \mu_0 + r} \frac{\kappa}{\kappa + r} \lambda_1 w_1 \left(\frac{\mu_0 - \mu_1}{\mu_1 + r} \right) > 0. \end{aligned}$$

The left-hand side of this inequality reflects the entry deterrence effect of R&D. For the monopolist, R&D serves as a barrier to entry from potential imitators (i.e. $\mu_0 - \mu_1 > 0$).

As shown in Figure 3, the two curves $\hat{p}(c)$ and $p^*(c)$ must eventually cross. This implies that *over*-experimentation occurs for high values of c. Given that $\hat{p}(c) < p^*(c)$ for low c, the thresholds cross only if

$$\lim_{c \to \infty} \hat{p}(c) = \frac{1}{\hat{k}_3} > \frac{1}{k_3^*} = \lim_{c \to \infty} p^*(c),$$

which holds since

$$k_3^* = \frac{\lambda_0 + \mu_0}{\lambda_0 + \mu_0 + r} > \frac{\lambda_0}{\lambda_0 + \mu_0 + r} = \hat{k}_3.$$

Since $\bar{c} < \bar{c}^*$, the crossing of the private and social belief thresholds arrives sufficiently fast, i.e. at some $c < \bar{c}$, such that over-experimentation emerges for high values of c.

The latter inequality represents the fact that the society benefits from the technological progress whether it comes from the incumbent or an entrant (i.e. the additional μ_0 term on the left-hand side). When experimenting, the planner therefore pays for it regardless of entry; in contrast, the firm incurs c only if sales arrive before imitators. This wedge strengthens the value of R&D to the society vis-à-vis the firm such that, at large c, the planner is willing to invest at beliefs where the firm would not.[18]

To sum up, at low values of c, or when the surplus from R&D is large, under-experimentation occurs because the greater marginal benefit of R&D to the firm, due to the entry-deterrence effect, is the dominant force. As the cost of R&D increases, or as the surplus from R&D decreases, such effects diminish and become eventually dominated by the negative impact of R&D, which falls disproportionately on the planner's incentives for experimentation because of the externalities from knowledge spillovers.

Remark 2. In the analysis above, we treated the informational content of entry differently across the two technological stages. In $s = 0$, entry brings an immediate arrival of sales if $\theta = 1$, while the sales process is unaffected by entry in $s = 1$. Our main result of this section nonetheless holds under alternative formulations of the planner's dynamic programming problem.

In Appendices B.1 and B.2, we consider the following two alternatives: (i) entry leads to immediate resolution of uncertainty in $s = 1$ as well as in $s = 0$, and (ii) entry has no impact on learning in $s = 0$ as well as in $s = 1$. In either case, the private and social belief thresholds diverge in such a way that both under- and over-experimentation emerge.

Remark 3. Suppose that the possibility of entry is shut down altogether but the monopolist earns $\pi_s = \rho_s w_s$, $\rho_s \in (0,1)$, from each arrival of sales in $s = 0, 1$. The planner maximizes the consumer surplus.

By setting $\mu_0 = \mu_1 = 0$ in the problem of both the firm and the planner, and replacing w_1 by $\rho_1 w_1$ in the firm's problem, we obtain

$$\hat{p}(c) = \frac{(\lambda_0 + r)c}{\frac{\kappa}{\kappa+r}\lambda_1 \rho_1 w_1 + \lambda_0 c}$$

and

$$p^*(c) = \frac{(\lambda_0 + r)c}{\frac{\kappa}{\kappa+r}\lambda_1 w_1 + \lambda_0 c}.$$

It is then straightforward to see that we must have $\hat{p}(c) > p^*(c)$ for all c, and hence, only over-experimentation. This demonstrates that knowledge spillovers are indeed responsible for the occurrence of both under- and over-experimentation in our model.

[18] This phenomenon is independent of the learning effect of entry. See Appendix B.2.

5 Patents

5.1 Overview

Our model admits two types of market failure: under- and over-experimentation. To address these inefficiencies, we investigate the role of patents.

Consider a two-tier patent system that consists of “regular” and “petty” patents. Regular patents protect innovations from imitators in the mature technological stage $s = 1$, while early-stage innovations in $s = 0$ are eligible for petty patents. As discussed further in Section 5.4 below, one way to interpret the distinction between regular and petty patents is differentiated “patentability” standards. That is, the innovator is subject to a less stringent set of standards for petty patent than regular patent.

We propose the following patent policy. For low-surplus innovations (i.e. $c > \gamma$), only regular patents are offered; for high-surplus innovations (i.e. $c < \gamma$), both petty and regular patents are offered. We first derive the required length of each patent as a function of c as well as the firm’s belief on market demand, p. We then examine the scope of the proposed policy when the regulator cannot elicit the firm’s privately held belief.

Another trade-off is introduced at this juncture. The patenting decision is voluntary, and if the firm chooses to sell the product without patent, it resorts to trade secret. Patenting reveals the content and feasibility of the technology, and this magnifies knowledge spillovers such that, upon expiry, the entry rate increases by a proportion $\alpha \geq 1$. That is, in each $s \in \{0, 1\}$, while trade secret maintains the entry rate at μ_s, patenting leads to a higher rate $\alpha\mu_s$.

5.2 Over-experimentation

We begin by considering the case of over-experimentation, as the result here will be used in addressing the more involved under-experimentation case below. When $c > \gamma$, we have $\hat{p}(c) > p^*(c)$, and the monopolist experiments at beliefs $p \in (p^*(c), \hat{p}(c))$ where the planner would invest in R&D instead.

The next result shows that socially optimal experimentation can be achieved via patents offered only in the mature technological stage $s = 1$. We derive the length of such a patent in terms of the belief on market demand as well as the cost of R&D. However, the scope of the policy is not limited by the knowledge of p, as explained in Remark 4.

Proposition 4. *Fix any $c \in (\gamma, \bar{c})$, where γ is given by Proposition 3, and any $p \in [p^* (c), \hat{p}(c))$. Then, there exists a regular patent of finite length $T_1(c, p)$, available in $s = 1$, with the following properties:*

- *In $s = 0$, the monopolist strictly prefers to invest in R&D at any $p' > p$, to experiment at any $p' < p$, and is indifferent at p;*
- *In $s = 1$, the monopolist strictly prefers to acquire the patent.*

Proof. See Appendix A. □

Note that the decision problem effectively ends once the firm engages in R&D. A regular patent offered in $s = 1$ with sufficiently lengthy protection allows the firm to appropriate greater profits from its scaled-up innovation and hence improves the continuation value associated with R&D. This encourages the firm to pursue R&D, which is similar to the standard effect of a patent. In our model, however, such effects are manifested in lowering the firm's belief threshold to align the private and social incentives for learning and to fasten the pace of innovation.

Figure 4 illustrates how the patent operates in this case. Fix any $c' \in (\gamma, \bar{c})$ and $p' \in [p^*(c'), \hat{p}(c'))$. Then, assuming that it gets taken up in $s = 1$ at the same belief, a regular patent of length $T_1(c', p')$ reduces the firm's threshold from $\hat{p}(c')$ precisely to p'. Thus, at belief p', the firm (weakly) prefers R&D to experimentation, just as the planner does. We show that the length $T_1(c', p')$ indeed provides the firm with strict incentives to obtain the patent upon completion of the R&D.

Note that uncertainty plays no role in the firm's patenting decision in $s = 1$; all that matters here is the trade-off between protection during the life of patent and the extra spillover effect upon expiry. Also, Proposition 4 is independent of the discount rate since there is no need to induce additional learning.

Remark 4. Suppose that p is unobservable. Then, we can simply implement regular patent with length $T_1(c, p^*(c))$, which has the effect of reducing the monopolist's belief threshold $\hat{p}(c)$ to the social threshold $p^*(c)$. Note, however, that $T_1(c, p^*(c)) > T_1(c, p)$ for all $p \in (p^*(c), \hat{p}(c))$. In other words, to accommodate asymmetric information on market belief, we need a patent longer than minimally necessary.

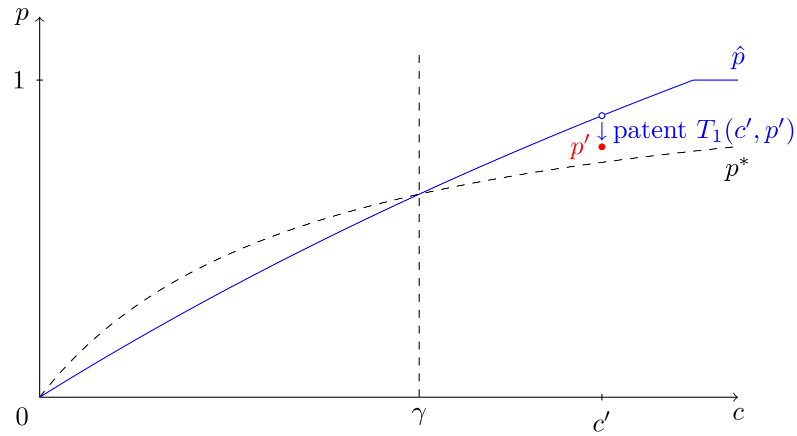

Figure 4: Correcting over-experimentation

Remark 5. Note that, for $\alpha > 1$,

$$\lim_{p \to \hat{p}(c)} T_1(c, p) = \frac{1}{r} \ln \left[\frac{\alpha(\mu_1 + r)}{\alpha\mu_1 + r} \right] > 0.$$ [19]

That is, even when the firm's threshold needs to be adjusted by an arbitrarily small amount, we need a patent length bounded away from zero to correct the inefficiency. This is due to the spillover effects of patenting that intensifies entry.

5.3 Under-experimentation

We now consider the case of under-experimentation. When $c < \gamma$, we have $\hat{p}(c) < p^*(c)$, and the monopolist invests in R&D at beliefs $p \in (\hat{p}(c), p^*(c))$, where the planner would experiment instead. The planner offers petty patent of length T_0 to innovations in $s = 0$, in addition to regular patent of length T_1 to innovations in $s = 1$.

As we saw in Section 5.2, regular patents have the effect of reducing the firm's threshold and would therefore exacerbate the divergence between the private and social incentives in the case of under-experimentation. However, even if regular patents are ruled out, the regulator may not be able to stop the firm from taking the same petty patent in both stages. After all, if an invention is patentable in $s = 0$, it should be patentable in $s = 1$. Furthermore, regular patents exist in practice to serve other functions not modeled explicitly here, especially to

[19] For the derivation, see (10) in Appendix A.

stimulate *ex ante* investment. For these reasons, we want to ensure that effective reward for early-stage innovations can be designed when regular patents are available at the same time.

Specifically, we ask whether, for given cost $c \in (0, \gamma)$ and belief $p \in (\hat{p}(c), p^*(c))$ as well as regular patent of length $T_1 \in [0, \infty)$, there exists a petty patent of finite length T_0 such that (i) the firm in stage $s = 0$ with belief p obtains the petty patent and experiments; and (ii) the firm prefers experimentation to R&D throughout and beyond the life of the patent, should no sales arrive.

At $p \in (\hat{p}(c), p^*(c))$, R&D is superior to experimentation under trade secret, and hence, the first condition is analogous to experimentation under the patent being preferred to R&D. This condition is non-trivial because patenting magnifies the spillover effects such that the entry rate jumps from μ_0 to $\alpha\mu_0$ upon expiry. The second condition implies a lower bound on T_0. The minimal length of patent, conditional on no sales, lets the belief drift down precisely to the firm's threshold at termination. Note, however, that this threshold differs from $\hat{p}(c)$ if we have either $T_1 > 0$ or $\alpha > 1$, or both.

In our model without patents, if the firm prefers experimentation over R&D at any given belief, it continues to do so if sales do not arrive, since the belief is then decreasing over time (see Figure 1). This remains true under a petty patent at least during its life.[20]

Lemma 3. *Fix any* $c \in (0, \gamma)$*, any* $p \in (\hat{p}(c), p^*(c)]$*, and regular patent of any length* $T_1 \in [0, \infty)$*. Consider a petty patent of length* $T_0 \geq 0$*. In* $s = 0$*, if the monopolist acquires the petty patent then it experiments throughout the life of the patent, conditional on no sales arriving.*

Proof. See Appendix A. □

Using this lemma, we establish a patent policy that induces a cutoff behavior from the monopolist, provided that the discount rate is sufficiently close to zero. The length of regular patent, T_1, or whether or not the firm obtains this patent in $s = 1$, is immaterial for the result.

Proposition 5. *Fix any* $c \in (0, \gamma)$ *and any* $p \in (\hat{p}(c), p^*(c)]$*. Also, fix regular patent of any length* $T_1 \in [0, \infty)$*. Then, there exists* $\bar{r} > 0$ *such that, for all* $r \in (0, \bar{r})$*, there exists a petty patent of finite length* $T_0(c, p, T_1)$*, available in* $s = 0$*, with the following properties:*

- *In* $s = 0$*, the monopolist strictly prefers to invest in R&D at any* $p' > p$*, to acquire the petty patent and experiment at any* $p' < p$*, and is indifferent at* p*.*

[20] We can in fact use the arguments for Lemma 3 more generally to set up the differential equation for the value function directly, as in Choi (1991). The methods in Section 3 are based on Keller and Rady (2015).

Proof. See Appendix A. □

Together with Lemma 3, Proposition 5 immediately implies the following.

Corollary 4. *As $r \to 0$, there exists a petty patent with sufficiently long life such that the monopolist with belief p acquires the patent and experiments throughout the continuation play, conditional on no sales arriving.*

The standard effect of a patent is to stimulate R&D by allowing the innovator to appropriate a larger share of the profits resulting from the invention. This is equivalent to lowering the monopolist's threshold $\hat{p}$ in our setup with over-experimentation. Proposition 5 highlights a potentially new role for patents. When the surplus from R&D is high, the firm invests in R&D at beliefs where it is instead socially optimal to gather information about demand. Introducing a petty patent to protect the early-stage innovation from imitators enables the firm to experiment in the market with greater confidence and slow down the pace of innovation. The society benefits from the reduced hazard of wasteful R&D activities.

To prove Proposition 5, we explicitly construct the value of experimentation under a fixed pair of patent lenghts, (T_0, T_1). This is a linear function in belief p, and admits the following cutoff property: there exists some $p(T_0)$ such that, for all $p \leq p(T_0)$, the monopolist prefers to take the patent, and vice versa. We then show that, for an under-experimenting monopolist, i.e. $p \in (\hat{p}(c), p^*(c)]$, there exists a finite length T_0 that indeed makes the firm indifferent between the patent and R&D, as long as the firm is sufficiently patient. While such T_0 does not guarantee the belief to drift down enough to a level upon termination where the learning incentive dominates the R&D incentive, we can always add more protection, which yields Corollary 4.

Note that petty patents cannot provide the desired incentives if r is too high, unlike Proposition 4 which does not depend on the discount rate. Here, the actual payoffs accrue only after the firm successfully learns good demand but the learning itself may take a long time.

The effects of petty patent are illustrated for the case of $\alpha = 1$ and $T_1 = 0$ in Figure 5. The left panel (Figure 5a) shows how the firm's belief evolves over time, while the right panel (Figure 5b) depicts the arguments in the (c, p) space. For simplicity, we let $T_0(c, p) := T_0(c, p, 0)$ for any c and p. Consider some $c' < \gamma$ and $p' \in (\hat{p}(c'), p^*(c'))$. The planner offers a patent only in $s = 0$. Assuming that the firm obtains the patent and experiments, absent any sales, the belief drifts down according to the curve $p(t)$. Here, $\widetilde{T}_0$ represents the minimal patent length required to achieve optimal experimentation.[21] Under the patent of length $T_0(c', p') \geq \widetilde{T}_0$, both the firm and the planner favor experimentation over R&D once the patent expires.

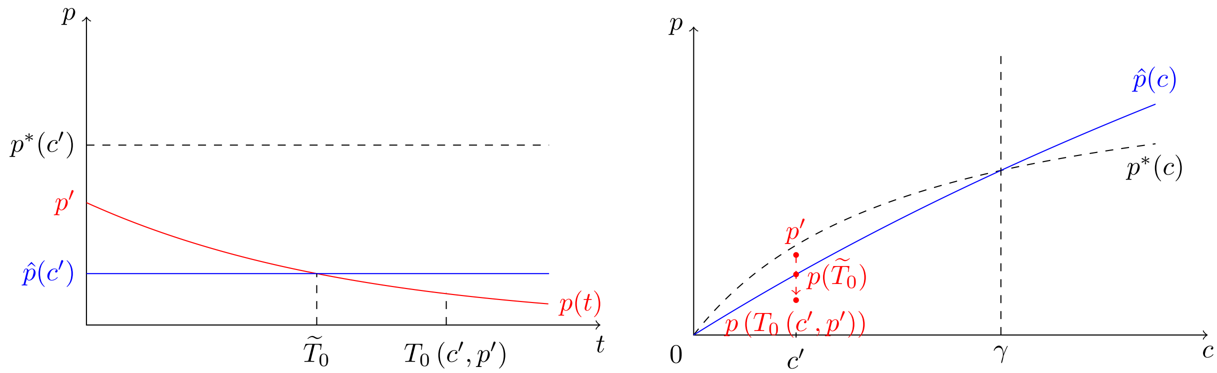

(a) Evolution of the belief over time (b) Evolution of the belief in the (c, p) plane

Figure 5: Correcting under-experimentation

Remark 6. Suppose that p is unobservable. Then, we can implement petty patent with length $T_0(c, p^*(c), T_1)$, which will be obtained by the monopolist if $p < p^*(c)$. Note the following. First, $T_0(c, p^*(c), T_1)$ may not be long enough for the belief to fall below the threshold at the end of patent life. If so, the effect of the patent would only be to delay R&D; to fully correct under-experimentation, we would want the monopolist to experiment at all $p < p^*(c)$. Second, such a patent is obtained not only by the under-experimenting monopolists (i.e. $p \in (\hat{p}(c), p^*(c))$) but also by those who would have experimented anyway (i.e. $p < \hat{p}(c)$). As discussed in Section 5.4 below, however, this latter issue does not have to raise a cause for concern in our learning context.

5.4 Discussion and Taking Stock

Summary When the surplus from R&D is low, the monopolist in our model over-experiments. This case corresponds to the well-known problem of under-investment in R&D due to positive

[21] If $\alpha > 1$ (or $T_1 > 0$), the post-patent belief threshold would fall below $\hat{p}(c')$. See part (iii) of Corollary 1.

externalities, and the inefficiency is corrected via regular patents available in $s = 1$ (Section 5.2). When the surplus from R&D is high, the threat of knowledge spillovers combined with the entry-deterrence effect of R&D generates under-experimentation and over-investment. Protecting early-stage innovations from imitators, in the form of petty patents available in $s = 0$, can restore the efficient level of experimentation in this case (Section 5.3). Note that offering only regular patents here would worsen the under-experimentation problem by raising the value of R&D.

Policy Implications We can interpret petty patents as imposing lower standards of patentability than regular patents. Indeed, outside the US, more than sixty countries, including Germany, Japan, and Korea, have been implementing two-tier patent systems with differentiated patentability standards. Second-tier patents in these systems, sometimes referred to as "utility models", typically require relaxed patentability standards, and are cheaper and quicker to acquire with shorter life of protection.[22]

Despite their global prevalence and on-going debates in the US (e.g. Janis, 1999; Sichelman, 2010), the economics literature is yet to provide a formal account of the role of petty patents. This paper proposes a theory from a learning perspective. For countries like the US, granting protection for minor inventions and less than fully commercialized products (e.g. beta versions) may improve welfare by stimulating experimentation and preventing socially wasteful R&D activities. For those already operating two-tier systems, our theory suggests restricting access to petty patents for only high-surplus innovations.

How can we distinguish high-surplus from low-surplus innovations to deliver the petty patents efficiently? One possibility is to broaden the contents of adjustment in the patentability standards. The main ingredients of patentability are novelty, non-obviousness, and utility. Among these, petty patents are usually associated with reduced burden on non-obviousness. By also considering the utility criterion, for instance, there could be a practical solution to the problem of screening high-surplus innovations.

Asymmetric Information In this paper, we showed that the scope of our results is not limited by regulatory knowledge of the firm's belief on market demand, although informational asymmetry may lead to longer patent life than minimally necessary to correct over-

[22]For more information on utility models, see https://www.wipo.int/patents/en/topics/utility_models.html.

experimentation (Remark 4), or only partial remedy for under-experimentation (Remark 6). Nonetheless, the benefits from introducing petty patents still depend on the surplus from R&D, which we assumed observable throughout.

A natural extension of our model would be to consider an incentive patent mechanism with unknown cost c. We not only want to screen high-surplus cases to award early-stage incentives, but the required length of patent depends on c whether we have high- or low-surplus innovation. As mentioned above, petty patents typically involve a host of terms that differ from regular patents in practice. Explaining the features of a two-tier patent system from a mechanism design perspective itself poses an interesting topic for future research even without learning dynamics. One potential challenge for screening in our learning setup however is that the gap between $\hat{p}$ and p^* responds non-monotonically to c (see Figure 3).

Other Trade-offs Our policy objective differs from the existing studies on sequential innovation. These papers deal with the trade-off between incentives for investment and the speed of knowledge accumulation. As such, with complete information, monetary incentives are superior to patents since they encourage investments without impeding follow-on innovations.[23] In our model, subsidizing or taxing early-stage innovation shifts the value of experimentation relative to R&D and hence provides an alternative policy instrument. But, R&D by entrants is identical to R&D by the incumbent, and this means that delayed arrival of entry does not necessarily erode welfare.[24] Transfer policy does not present any intrinsic advantage over patent policy in our learning context.

6 Conclusion

This paper considered a monopolist's incentives for learning about uncertain demand in the context of innovation and market competition. The monopolist chooses to invest in costly R&D to scale up the value of its innovation and to improve the prospect of a quiet life, or experiment by selling an early version of the technology at the expense of inviting imitators. We characterized the firm's optimal path of appropriation and analyzed the impact of various

[23]With incomplete information, patents may be a better policy option as they offer screening devices in addition to transfers (e.g. Scotchmer, 1999; Cornelli and Schankerman, 1999).

[24]This claim is true in the case with no learning effects (Remark 2 and Appendix B.2). Even when entry facilitates learning, note that patenting itself creates spillover effects which may compensate for delayed entry.

market characteristics, including the price and speed of learning as well as the intensity of knowledge spillovers.

Socially optimal experimentation takes account of the social returns to knowledge spillovers and therefore diverges from the profit-maximizing behavior. The novel theoretical insight that emerges from our analysis is that the multiple facets of knowledge spillovers can generate distinct patterns of market failure: both under- and over-experimentation are possible in our model. While the standard effect of knowledge spillovers is to strengthen the social value of R&D and cause under-investment, we identify a contrarian force in the entry-deterrence effect of R&D. As a result, under-investment in R&D, or over-experimentation, occurs only when the surplus from investment is low; otherwise, the firm performs R&D at beliefs where it is socially beneficial to experiment in the market.

The key property of innovation as a driver of increasing returns and long-run growth is its public good nature (e.g. Romer, 1990), and patents have long been regarded as a major source of balancing these *ex post* externalities with *ex ante* incentives for investment. In this paper, we emphasized an additional trade-off between knowledge spillovers and learning, and proposed a new role for patents. A patent policy that differentiates the patentability standard can control the pace of innovation and restore the socially optimal level of experimentation in our model.

For low-surplus innovations, regular patents offered to mature products suffice to discourage experimentation and stimulate investment, similarly to the classic case of under-investment in R&D. For high-surplus innovations, our analysis presents a case for reducing patentability. Protecting minor inventions may bring welfare gains by allowing innovators to experiment in the market early with greater confidence. Our analysis provides a theoretical foundation for two-tier patent system with contingent exercise of petty patents based on the potential value of the innovation.

Dropbox vs. Google Glass What can we then make of the contrasting episodes of Dropbox and Google Glass mentioned in Introduction? Dropbox learned from an instructive video that revealed only the basic features of its technology. In contrast, when a marketable prototype of Google Glass was released, it was already a mature product and sold as such, at the full retail price.[25] Our theory suggests that this difference in corporate strategy could be

[25]Prior to the prototype, "Explorer Edition", Google did offer demonstrations of the concept but it is not clear whether consumers were able to register their demand at the time. Dropbox had a beta waiting list

attributed to the size of the surplus at stake. From the society's standpoint, Dropbox may not have needed the video and Google was perhaps too afraid of leaking its know-how.

Patenting strategies of the two firms further corroborate the incentives that we model. Dropbox posted the video and launched the product in 2008, but its first patent, "Network folder synchronization", was filed for application later in 2009.[26] Google's patent application for "Wearable device with input and output structures" was made in August 2011, several years before the release of the prototype (April 2013) and the open beta (May 2014). The patent, however, was only granted in March 2016; moreover, the patent application was disclosed in February 2013, just two months prior to the prototype.[27]

While it is possible to reconcile the timeline of Dropbox with our analysis of Section 5.2, the case of Google Glass points to a potential conflict between the existing patent system and desirable pace of innovation. The long pendency may have forced Google to accelerate its efforts towards technological maturity. Petty patents with relaxed patentability requirement and shorter decision lag might have saved the company from wasting millions on developing a product that turned out to have only little commercial applicability.[28]

alongside the video.

[26]Dropbox applied for the patent provisionally in August 2009, with the formal application following a year later. The grant date is in September 2014. See https://patents.google.com/patent/US8825597B1.

[27]See https://patents.google.com/patent/US9285592B2.

[28]Google Glass was intended for the general consumer. After its failure, the company moved into new markets, targeting specific industrial users with "Google Glass Enterprise Edition". For more information, see https://en.wikipedia.org/wiki/Google_Glass.

Appendix

A Omitted Proofs

Proof of Corollary 1

Recall that $\hat{p}$ is given by

$$\hat{p} = \frac{(\lambda_0 + \mu_0 + r)c}{\frac{\kappa}{\kappa+r}\frac{\mu_0+r}{\mu_1+r}\lambda_1 w_1 + \lambda_0 c}.$$

It is straightforward to see that $\hat{p}$ is increasing with c and μ_1, and decreasing with w_1, κ, and λ_1.

The partial derivative with respect to λ_0 is given by

$$\frac{\partial \hat{p}}{\partial \lambda_0} = \frac{c\left[\frac{\kappa}{\kappa+r}\frac{\mu_0+r}{\mu_1+r}\lambda_1 w_1 - (\mu_0 + r)c\right]}{\left[\frac{\kappa}{\kappa+r}\frac{\mu_0+r}{\mu_1+r}\lambda_1 w_1 + \lambda_0 c\right]^2}.$$

The numerator, and hence the derivative, are positive if

$$c < \frac{\lambda_1 \kappa}{(\kappa + r)(\mu_1 + r)},$$

which holds for

$$c < \bar{c} = \frac{\lambda_1 \kappa}{(\kappa + r)(\mu_1 + r)} w_1.$$

Finally, the partial derivative with respect to μ_0 is given by

$$\frac{\partial \hat{p}}{\partial \mu_0} = -\frac{c\left[\frac{\lambda_0 \lambda_1 \kappa}{(\kappa+r)(\mu_1+r)} w_1 - \lambda_0 c\right]}{\left[\frac{\kappa}{\kappa+r}\frac{\mu_0+r}{\mu_1+r}\lambda_1 w_1 + \lambda_0 c\right]^2}.$$

This is negative also for $c < \bar{c}$.

Proof of Corollary 2

Recall that p^* is given by

$$p^* = \frac{(\lambda_0 + \mu_0 + r)c}{\frac{\kappa}{\kappa+r}\lambda_1 w_1 + (\lambda_0 + \mu_0)c}.$$

It is straightforward to see that p^* increasing with c, decreasing with w_1, λ_1, and κ, and does not depend on μ_1.

We now look at how p^* changes with λ_0 and μ_0. Note that λ_0 and μ_0 are interchangeable in the expression of p^*, so they have the same effect. We have

$$\frac{\partial p^*}{\partial \lambda_0} = \frac{\left[\frac{\kappa}{\kappa+r}\lambda_1 w_1 - rc\right] c}{\left[\frac{\kappa}{\kappa+r}\lambda_1 w_1 + (\lambda_0 + \mu_0)c\right]^2},$$

which has a positive numerator for $c < \bar{c}^*$.

Proof of Proposition 4

Fix $c \in (\gamma, \bar{c})$ and $p \in [p^*(c), \hat{p}(c))$. The proof proceeds in three steps. First, we find the minimal patent length such that the firm is willing to patent in $s = 1$ at belief p. Then, we find the length of patent (assuming that it is obtained in $s = 1$ at belief p) that makes the firm indifferent between R&D and experimentation in $s = 0$ at belief p. Finally, we verify that the latter length is greater than the former.

Step 1 Without a patent, or under trade secret, recall from Lemma 1 that the value after having innovated at belief p is

$$U_1(p) = \frac{\lambda_1 p}{\mu_1 + r} w_1.$$

Consider now a situation in which a regular patent has been awarded with remaining life T_1. Let $V_1(p, T_1)$ denote the value under the patent. Again, $V_1(p, T_1)$ must be linear in p, i.e. $V_1(p, T_1) = V_1(1, T_1)p$.

While the patent shuts down entry during its life, the rate of entry jumps to $\alpha\mu_1$ upon expiry. Nonetheless, except for the entry rate, the firm's post-patent incentives are the same as in Section 3. Thus, we must have

$$V_1(1, 0) = \frac{\lambda_1}{\alpha\mu_1 + r} w_1.$$

Now, consider the value when the patent is still valid, that is, when $T_1 > 0$. During a small interval of time dT_1, sales occur with probability $1 - e^{-\lambda_1 dT_1}$, yielding instantaneous payoff of w_1 and continuation value $V_1(1, T_1 - dT_1)$. We therefore have

$$V_1(1, T_1) = (1 - e^{-\lambda_1 dT_1})w_1 + e^{-r dT_1} V_1(1, T_1 - dT_1).$$

Using a first order approximation, subtracting $V_1(1, T_1 - dT_1)$ from both sides, dividing by dT_1, and letting $dT_1 \to 0$, we then obtain

$$\frac{dV_1(1,T_1)}{dT_1} = \lambda_1 w_1 - rV_1(1,T_1),$$

which, together with the boundary condition $V_1(1,0) = \lambda_1 w_1/(\alpha\mu_1 + r)$, gives

$$V_1(1,T_1) = \frac{\lambda_1 w_1}{r}\left(1 - \frac{\alpha\mu_1}{\alpha\mu_1 + r}e^{-rT_1}\right). \tag{9}$$

By the linearity of $V_1(p, T_1)$, it follows that

$$V_1(p,T_1) = \left[\frac{\lambda_1}{r}w_1 - \frac{\alpha\mu_1\lambda_1}{(\alpha\mu_1 + r)r}w_1 e^{-rT_1}\right]p.$$

We can then see that there is a unique length $\widetilde{T}_1$ such that $V_1(p, \widetilde{T}_1) = U_1(p)$, or

$$\widetilde{T}_1 = \frac{1}{r}\ln\left[\frac{\alpha(\mu_1 + r)}{\alpha\mu_1 + r}\right].$$

Step 2 Fix a patent of length $T_1 > \widetilde{T}_1$ such that the firm strictly prefers to patent once R&D is finished in $s = 1$, and consider the incentives for performing R&D in $s = 0$.

Since the firm would patent the innovation in $s = 1$, $V_1(1,T_1) > U_1(T_1)$. We can then reproduce the analysis of Section 3, replacing $U_1(1)$ by $V_1(1,T_1)$ in (2) and in Proposition 1. This allows us to find the new threshold

$$\begin{aligned}\hat{p}_1(c,T_1) &= \frac{(\lambda_0 + \mu_0 + r)c}{\frac{\kappa}{\kappa+r}(\mu_0 + r)V_1(1,T_1) - \lambda_0 w_0 + \lambda_0 c} \\ &= \frac{(\lambda_0 + \mu_0 + r)c}{\frac{\kappa}{\kappa+r}(\mu_0 + r)\frac{\lambda_1}{r}w_1\left[1 - \frac{\alpha\mu_1}{(\alpha\mu_1+r)}e^{-rT_1}\right] + \lambda_0 c},\end{aligned}$$

where $V_1(1,T_1)$ is given by (9).

Then, by solving $\hat{p}_1(c, T_1(c,p)) = p$, we obtain

$$T_1(c,p) = \frac{1}{r}\ln\left[\frac{\frac{\kappa}{\kappa+r}\frac{(\mu_0+r)\alpha\mu_1}{\alpha\mu_1+r}\frac{\lambda_1}{r}w_1 p}{\left(\frac{\kappa}{\kappa+r}\frac{\mu_0+r}{r}\lambda_1 w_1 + \lambda_0 c\right)p - (\lambda_0 + \mu_0 + r)c}\right]. \tag{10}$$

Note that, since the threshold is now such that $\hat{p}_1(c, T_1(c,p)) = p$, the monopolist with any belief $p' > p$ ($p' < p$) prefers to invest in R&D (experiment).

Step 3 Finally, we need to check that $T_1(c, p) \geq \widetilde{T}_1$ for all $p \in (p^*(c), \hat{p}(c))$. This inequality reduces to $p \leq \hat{p}(c)$, and since $p < \hat{p}(c)$ by assumption, we in fact have $T_1(c, p) > \widetilde{T}_1$.

Proof of Lemma 3

Recall from the proof of Proposition 4 that the value in $s = 1$, given T_1 and belief $p(t)$, is linear in $p(t)$, i.e.

$$V_1(p(t), T_1) = V_1(1, T_1)p(t),$$

where $V_1(1, T_1)$ is given by (9).

Suppose that the monopolist with belief $p(t)$ in $s = 0$ has obtained the petty patent with length T_0. Consider a small interval of time dt and the choice between switching to R&D immediately and experimenting for dt, under the patent protection, followed by switching to R&D.

The value of switching to R&D immediately is given by

$$-c + \frac{\kappa}{\kappa + r} V_1(1, T_1)p(t), \tag{11}$$

while the value of further experimentation for dt followed by R&D is given by

$$p(t)\lambda_0 dt \left[-c + \frac{\kappa}{\kappa + r} V_1(1, T_1)\right] + [1 - p(t)\lambda_0 dt - r dt] \left[-c + \frac{\kappa}{\kappa + r} V_1(1, T_1)p(t + dt)\right]. \tag{12}$$

Using the approximation $p(t + dt) \simeq p(t) - \lambda_0 p(t)(1 - p(t))dt$, and letting dt go to zero, (12) exceeds (11) if

$$-c + \frac{\kappa r}{\kappa + r} V_1(1, T_1)p(t) < 0.$$

Note that, conditional on no sales, $p(t)$ is decreasing in t. This implies that if the monopolist prefers the patent and experimentation over R&D now then it must do so throughout the life of the patent.

Proof of Proposition 5

Fix $c \in (0, \gamma)$ and regular patent of length $T_1 \in [0, \infty)$. The proof proceeds in two steps. First, we show that, for a given length T_0 of petty patent, the policy of the monopolist is in cutoff: there is a critical type $\hat{p}(T_0) \in (0, 1)$ such that, for all $p < \hat{p}(T_0)$, the firm obtains the patent and experiments, and for all $p > \hat{p}(T_0)$, the firm invests in R&D. Second, we show that, for any belief $p \in (\hat{p}(c), p^*(c)]$, we can find some finite T_0 that makes the monopolist indifferent between R&D and experimentation with the corresponding petty patent.

Step 1 Fix T_0. Also, fix any $p \in [0,1]$. The value of R&D is given by

$$-c + \frac{\kappa}{\kappa + r} V_1(1, T_1)p,$$

where $V_1(1, T_1)$ is given by (9). This is affine in p with a negative intercept.

We now construct the value from taking the petty patent with length T_0 and experimenting at p. By Lemma 3, we know that if the patent is taken then, absent any sales, the monopolist will experiment until the patent expires. Let $p(t)$ and $V_0^{pat}(t)$ denote the corresponding belief and value, respectively, after time $t \in [0, T_0]$ has passed without sales since the beginning of the patent. Note that $p := p(0)$.

For notational simplicity, let $K := -c + \frac{\kappa}{\kappa+r} V_1(1, T_1)$, and consider a small interval of time dt at time t. We have the following:

- With probability $p(t)\lambda_0 dt$, there is a sale and the payoff is K;
- With the complementary probability, there is no sale and the continuation value is $V_0^{pat}(t + dt)$.

Thus, we have

$$V_0^{pat}(t) = p(t)\lambda_0 dtK + (1 - p(t)\lambda_0 dt - rdt)V_0^{pat}(t + dt).$$

Subtracting $V_0^{pat}(t + dt)$ from both sides, dividing by dt, and letting $dt \to 0$, it then follows that V_0^{pat} must solve the following differential equation:

$$-V'(t) + (p(t)\lambda_0 + r)V(t) = p(t)\lambda_0 K. \tag{13}$$

To solve (13), we must use the value upon the patent's expiry, $V_0^{pat}(T_0)$, as the terminal condition. At $p(T_0)$, there are two cases to consider: (i) the monopolist prefers experimentation and (ii) the monopolist prefers R&D.

To find $V_0^{pat}(T_0)$, we can replicate the analysis of Section 3, with two modifications. First, once the patent expires in $s = 0$, the entry rate jumps from μ_0 to $\alpha\mu_0$; second, the value in $s = 1$ is given by $V_1(p, T_1)$ instead of $U_1(p)$.

Therefore, the value upon expiry is either

$$V_0^{pat}(T_0) = \frac{\lambda_0}{\lambda_0 + \alpha\mu_0 + r} \left[-c + \frac{\kappa}{\kappa + r} V_1(1, T_1)\right] p(T_0) \tag{14}$$

if experimentation continues to be preferred, or

$$V_0^{pat}(T_0) = -c + \frac{\kappa}{\kappa + r} V_1(1, T_1) p(T_0) \tag{15}$$

if the monopolist switches to R&D. Note that the monopolist is indifferent between experimenting and doing R&D upon expiry at belief

$$\hat{p}_{pat}(c) = \frac{(\lambda_0 + \alpha\mu_0 + r)c}{\frac{\kappa}{\kappa+r}(\alpha\mu_0 + r)V_1(1, T_1) + \lambda_0 c} \tag{16}$$

We consider the two cases in turn.

Case 1: Experimentation is preferred upon expiry.

Let us first consider the case in which experimentation is preferred at the end of patent life. To solve the problem in this case, assume that $V_0^{pat}(t) = f(t)p(t)$.

Using the law of motion for beliefs, we have

$$\frac{dV_0^{pat}(t)}{dt} = f'(t)p(t) + f(t)p'(t) = f'(t)p(t) - f(t)\lambda_0 p(t)(1 - p(t)),$$

and (13) becomes

$$-f'(t) + (\lambda_0 + r)f(t) = \lambda_0 K,$$

which has a solution of the form

$$f(t) = \frac{\lambda_0}{\lambda_0 + r} K + Ae^{(\lambda_0 + r)t},$$

where A is a constant.

Plugging the above back into $V_0^{pat}(t) = f(t)p(t)$, and substituting for K, we obtain

$$V_0^{pat}(t) = \left[\frac{\lambda_0}{\lambda_0 + r}\left[-c + \frac{\kappa}{\kappa + r}V_1(1, T_1)\right] + Ae^{(\lambda_0 + r)t}\right] p(t).$$

We can then use the terminal condition (14) to determine A:

$$A = -\frac{\lambda_0 \alpha\mu_0}{(\lambda_0 + r)(\lambda_0 + \alpha\mu_0 + r)}\left[-c + \frac{\kappa}{\kappa + r}V_1(1, T_1)\right] e^{-(\lambda_0 + r)T_0}. \tag{17}$$

Note that A is independent of p, as the belief cancels out.

The value of experimenting at the beginning of the patent is thus given by

$$V_0^{pat}(0) = \left[\frac{\lambda_0}{\lambda_0 + r}\left[-c + \frac{\kappa}{\kappa + r}V_1(1, T_1)\right] + A\right]p, \tag{18}$$

which is linear in p.

Case 2: R&D is preferred upon expiry.

We next consider the second case, in which the monopolist prefers R&D at the end of patent life. First, note that we have the following particular solution for (13)

$$V(t) = \frac{\lambda_0}{\lambda_0 + r}Kp(t).$$

Thus, we search for a solution of the form

$$V_0^{pat}(t) = V(t) + v(t),$$

where $v(t)$ is the complementary function, that is, the solution to the homogenous differential equation

$$-v'(t) + [\lambda_0 p(t) + r]\, v(t) = 0.$$

By re-arranging and integrating, it then follows that

$$\int_0^t \frac{v'(s)}{v(s)}ds = \int_0^t [\lambda_0 p(s) + r]\, ds.$$

Recall from the law of motion of beliefs that $p'(s) = -\lambda_0 p(s)(1 - p(s))$, so that

$$\begin{aligned}\int_0^t \lambda_0 p(s) ds &= \int_0^t -\frac{p'(s)}{1 - p(s)}ds \\ &= \ln\left[1 - p(s)\right]\Big|_0^t \\ &= \ln\frac{1 - p(t)}{1 - p}.\end{aligned}$$

Therefore, we have

$$\int_0^t \frac{v'(s)}{v(s)}ds = \ln\frac{1 - p(t)}{1 - p} + rt,$$

so that

$$v(t) = Be^{rt}\frac{1 - p(t)}{1 - p},$$

where B is a constant.

Note that we can re-write

$$\frac{1-p(t)}{1-p} = \frac{1}{(1-p)+pe^{-\lambda_0 t}} = e^{\lambda_0 t}\frac{p(t)}{p},$$

implying

$$V_0^{pat}(t) = \frac{\lambda_0}{\lambda_0+r}Kp(t) + Be^{(\lambda_0+r)t}\frac{p(t)}{p}.$$

Using the terminal condition (15), we find the constant

$$\begin{aligned} B &= e^{-(\lambda_0+r)T_0}\left[-c\frac{p}{p(T_0)} + \frac{\lambda_0}{\lambda_0+r}cp + \frac{r}{\lambda_0+r}\frac{\kappa}{\kappa+r}V_1(1,T_1)p\right] \\ &= e^{-(\lambda_0+r)T_0}\left[-c\left[p+(1-p)e^{\lambda_0 T_0}\right] + \frac{\lambda_0}{\lambda_0+r}cp + \frac{r}{\lambda_0+r}\frac{\kappa}{\kappa+r}V_1(1,T_1)p\right], \end{aligned}$$

where we use the fact that $p/p(T_0) = p+(1-p)e^{\lambda_0 T_0}$. Thus, B is affine in p, and we write it as $B(p)$.

We can then plug the constant back into $V_0^{pat}(0)$ to get the value at the start of the patent, which is given by

$$V_0^{pat}(0) = \frac{\lambda_0}{\lambda_0+r}Kp + B(p).$$

Since B is affine in p, this is also affine in p.

From Case 1 and Case 2, we conclude that the initial value of the patent is either linear or affine in p, as is the value of R&D. Also, it is straightforward to see that, when $p=0$,

$$V_0^{pat}(0) = 0 > -c$$

while, if $p=1$,

$$V_0^{pat}(0) = \left[\frac{\lambda_0}{\lambda_0+r} + e^{-(\lambda_0+r)T_0}\frac{r}{\lambda_0+r}\right]\left[-c+\frac{\kappa}{\kappa+r}V_1(1,T_1)\right] < -c+\frac{\kappa}{\kappa+r}V_1(1,T_1),$$

where the right-hand side of the two inequalities corresponds to the value of R&D. Then, the intermediate value theorem implies that there is a unique threshold $\hat{p}(T_0) \in (0,1)$ such that the monopolist takes the patent and experiments if and only its belief is below $\hat{p}(T_0)$.

Step 2 If no sales arrive during the life of the patent, the belief drifts down to $p(T_0) = p/(p + (1-p)e^{\lambda_0 T_0})) \to 0$ as $T_0 \to \infty$. Thus, when T_0 is sufficiently large, the belief will be below the threshold $\hat{p}_{pat}(c)$ in (16), and the monopolist will prefer experimenting upon expiry. Moreover, the value of the patent will be given by (18) in Case 1 above.

Note that the value of A given in (17) goes to zero as $T_0 \to \infty$. Hence, the value of the patent converges to

$$[-c + V_1(1, T_1)]\, p$$

as $T_0 \to \infty$ and $r \to 0$. This is greater than the value of R&D, which is

$$-c + V_1(1, T_1)p$$

as $r \to 0$. While the expected benefits are the same, the patent allows to avoid the the cost of doing R&D when demand is bad.

When $T_0 = 0$, that is, when there is no patent, we know that the monopolist prefers to do R&D right away, as we are in the case where $p \in (\hat{p}(c), p^*(c)]$. Thus, there exists a finite value of $T_0 := T_0(c, p, T_1)$ that makes the firm with the given belief p indifferent between investing in R&D and experimenting under patent protection, and from the previous step, we know that the firm prefers R&D if $p' > p$ and experimentation if $p' < p$.

B Alternative Formulations of Dynamic Programming

B.1 Entry reveals demand in both $s = 0$ and $s = 1$

Suppose that entry generates immediate resolution of uncertainty in $s = 1$ as well as in $s = 0$. In this case, the planner's dynamic programming problem of Section 4.1 can be modified by replacing the arrival rate λ_1 with $\lambda_1 + \mu_1$, which implies

$$W_1(1) = \frac{\lambda_1 + \mu_1}{r} w_1.$$

All other aspects of the planner's problem remain identical.

It is then straightforward to re-calculate the social belief threshold:

$$\begin{aligned} p^*(c) &= \frac{(\lambda_0 + \mu_0 + r)c}{\frac{\kappa}{\kappa + r} r W_1(1) + (\lambda_0 + \mu_0)c} \\ &= \frac{(\lambda_0 + \mu_0 + r)c}{\frac{\kappa}{\kappa + r}(\lambda_1 + \mu_1)w_1 + (\lambda_0 + \mu_0)c}. \end{aligned}$$

Recall from Section 4.2 that, for the private and social belief thresholds to cross, three conditions need to be satisfied:

(i) $\hat{p}(0) = p^*(0) = 0$;

(ii) $\hat{p}'(0) < p^{*\prime}(0)$;

(iii) $\lim_{c\to\infty} \hat{p}(c) > \lim_{c\to\infty} p^*(c)$.

Conditions (i) and (iii) continue to hold, and condition (ii) now becomes

$$\frac{\kappa}{\kappa+r}\frac{1}{\mu_1+r}\left[\lambda_1(\mu_0-\mu_1)-\mu_1(\mu_1+r)\right]w_1 > 0,$$

which holds if the term inside brackets on the left-hand side is positive, or if

$$\frac{\mu_1}{\mu_0} < \frac{\lambda_1}{\lambda_1+\mu_1+r}. \tag{19}$$

If (19) is satisfied, the thresholds intersect at

$$\gamma = \frac{\kappa\lambda_1(\mu_0-\mu_1)}{\mu_0(\kappa+r)(\mu_1+r)}w_1 - \frac{\kappa\mu_1}{\mu_0(\kappa+r)}w_1.$$

B.2 Entry is uninformative in both $s = 0$ and $s = 1$

This case shuts down the learning effect of entry altogether. We can simply set $\mu_0 = \mu_1 = 0$ in the planner's problem and obtain

$$p^*(c) = \frac{(\lambda_0+r)c}{\frac{\kappa}{\kappa+r}\lambda_1 w_1 + \lambda_0 c}.$$

Again, we check the conditions for the crossing of the private and social thresholds:

(i) $\hat{p}(0) = p^*(0) = 0$;

(ii) $\hat{p}'(0) < p^{*\prime}(0)$;

(iii) $\lim_{c\to\infty} \hat{p}(c) > \lim_{c\to\infty} p^*(c)$.

Conditions (i) and (iii) remain valid, while condition (ii) now becomes

$$\frac{\kappa\lambda_1 w_1}{(\lambda_0 + r)(\lambda_0 + \mu_0 + r)(\kappa + r)(\mu_1 + r)} \left[\lambda_0(\mu_0 - \mu_1) - \mu_1(\mu_0 + r)\right] > 0,$$

which holds if

$$\frac{\mu_1}{\mu_0} < \frac{\lambda_0}{\lambda_1 + \mu_1 + r}. \tag{20}$$

If (20) is satisfied, the thresholds intersect at

$$\gamma = \frac{\kappa\lambda_1 w_1}{\lambda_0\mu_0(\kappa + r)(\mu_1 + r)} \left[\lambda_0(\mu_0 - \mu_1) - \mu_1(\mu_0 + r)\right].$$

References

Acemoglu, D., K. Bimpikis, and A. Ozdaglar (2011). Experimentation, patents, and innovation. *American Economic Journal: Microeconomics 3*, 37–77.

Aghion, P., R. Blundell, R. Griffith, P. Howitt, and S. Prantl (2009). The effects of entry on incumbent innovation and productivity. *Review of Economics and Statistics 91*, 20–32.

Akcigit, U. and Q. Liu (2015). The role of information in innovation and competition. *Journal of the European Economic Association 14*, 828–870.

Arrow, K. (1962). Economic welfare and the allocation of resources. In *The rate and direction of inventive activity: Economic and social factors*, pp. 609–626. Princeton University Press.

Arundel, A. (2001). The relative effectiveness of patents and secrecy for appropriation. *Research policy 30*, 611–624.

Bergemann, D. and U. Hege (1998). Venture capital financing, moral hazard, and learning. *Journal of Banking and Finance 22*, 703–735.

Bergemann, D. and U. Hege (2005). The financing of innovation: Learning and stopping. *RAND Journal of Economics 36*, 719–752.

Besanko, D. and J. Wu (2013). The impact of market structure and learning on the tradeoff between R&D competition and cooperation. *Journal of Industrial Economics 61*, 166–201.

Bolton, P. and C. Harris (1999). Strategic experimentation. *Econometrica 67*, 349–374.

Bonatti, A. and J. Hörner (2011). Collaborating. *American Economic Review 101*, 632–663.

Bonatti, A. and J. Hörner (2017). Career concerns with exponential learning. *Theoretical Economics 12*, 425–475.

Camargo, B. (2007). Good news and bad news in two-armed bandits. *Journal of Economic Theory 135*, 558–566.

Choi, J. P. (1991). Dynamic R&D competition under "hazard rate" uncertainty. *RAND Journal of Economics 22*, 596–610.

Cohen, W. M., R. R. Nelson, and J. P. Walsh (2000). Protecting their intellectual assets: Appropriability conditions and why us manufacturing firms patent (or not). National Bureau of Economic Research, w7552.

Cornelli, F. and M. Schankerman (1999). Patent renewals and R&D incentives. *RAND Journal of Economics 30*, 197–213.

Dasgupta, P. and J. Stiglitz (1980). Uncertainty, industrial structure, and the speed of R&D. *Bell Journal of Economics 11*, 1–28.

Denicoló, V. (2000). Two-stage patent races and patent policy. *RAND Journal of Economics 31*, 488–501.

Gilbert, R. J. and D. M. Newbery (1982). Preemptive patenting and the persistence of monopoly. *American Economic Review 72*, 514–526.

Green, J. R. and S. Scotchmer (1995). On the division of profit in sequential innovation. *RAND Journal of Economics 26*, 20–33.

Guo, Y. (2016). Dynamic delegation of experimentation. *American Economic Review 106*, 1969–2008.

Halac, M., N. Kartik, and Q. Liu (2016). Optimal contracts for experimentation. *Review of Economic Studies 83*, 1040–1091.

Halac, M., N. Kartik, and Q. Liu (2017). Contests for experimentation. *Journal of Political Economy 125*, 1523–1569.

Hayek, F. A. (1948). *Individualism and economic order*. University of Chicago Press.

Hicks, J. R. (1935). Annual survey of economic theory: the theory of monopoly. *Econometrica 3*, 1–20.

Hopenhayn, H., G. Llobet, and M. Mitchell (2006). Rewarding sequential firms: Prizes, patents, and buyouts. *Journal of Political Economy 114*, 1041–1068.

Hörner, J. and L. Samuelson (2013). Incentives for experimenting agents. *RAND Journal of Economics 44*, 632–663.

Jaffe, A. B., M. Trajtenberg, and R. Henderson (1993). Geographic localization of knowledge spillovers as evidenced by patent citations. *Quarterly Journal of Economics 108*, 577–598.

Janis, M. D. (1999). Second tier patent protection. *Harvard International Law Journal 40*, 151–219.

Keller, G. and S. Rady (2010). Strategic experimentation with poisson bandits. *Theoretical Economics 5*, 275–311.

Keller, G. and S. Rady (2015). Breakdowns. *Theoretical Economics 10*, 175–202.

Keller, G., S. Rady, and M. Cripps (2005). Strategic experimentation with exponential bandits. *Econometrica 73*, 39–68.

Kerr, W. R., R. Nanda, and M. Rhodes-Kropf (2014). Entrepreneurship as experimentation. *Journal of Economic Perspectives 28*, 25–48.

Klein, N. and S. Rady (2011). Negatively correlated bandits. *Review of Economic Studies 78*, 693–732.

Loury, G. C. (1979). Market structure and innovation. *Quarterly Journal of Economics 93*, 395–410.

Lundvall, B. (1988). Innovation as an interactive process: from user-producer interaction to the national system of innovation. In *Technical change and economic theory*, pp. 349–369. Pinter.

Malueg, D. A. and S. O. Tsutsui (1997). Dynamic R&D competition with learning. *RAND Journal of Economics 28*, 751–772.

Manso, G. (2011). Motivating innovation. *Journal of Finance 66*, 1823–1860.

Manso, G. (2016). Experimentation and the returns to entrepreneurship. *Review of Financial Studies 29*, 2319–2340.

Murto, P. and J. Välimäki (2011). Learning and information aggregation in an exit game. *Review of Economic Studies 78*, 1426–1461.

Nelson, R. R. (1977). *The moon and the ghetto.* Norton.

O'Donoghue, T. (1998). A patentability requirement for sequential innovation. *RAND Journal of Economics 29*, 654–679.

Parra, A. (2019). Sequential innovation, patent policy, and the dynamics of the replacement effect. *RAND Journal of Economics 50*, 568–590.

Reinganum, J. F. (1983). Uncertain innovation and the persistence of monopoly. *American Economic Review 73*, 741–748.

Romer, P. M. (1990). Endogenous technological change. *Journal of Political Economy 98*, S71–S102.

Rosenberg, N. (1994). *Exploring the black box: Technology, economics, and history.* Cambridge University Press.

Scotchmer, S. (1999). On the optimality of the patent renewal system. *RAND Journal of Economics 30*, 181–196.

Sichelman, T. (2010). Commercializing patents. *Stanford Law Review 62*, 341–412.

Spence, M. A. (1977). Entry, capacity, investment, and oligopolistic pricing. *Bell Journal of Economics 8*, 534–544.

Strulovici, B. (2010). Learning while voting: Determinants of collective experimentation. *Econometrica 78*, 933–971.

Von Hippel, E. (1988). *The sources of innovation.* Oxford University Press.

Weeds, H. (2002). Strategic delay in a real options model of R&D competition. *Review of Economic Studies 69*, 729–747.