Scarcity of Ideas and R&D Options:
Use it, Lose it, or Bank it[1]

Nisvan Erkal
University of Melbourne

Suzanne Scotchmer
University of California, Berkeley and NBER

This version: November 27, 2010

Abstract

We investigate optimal rewards in an R&D model where substitute ideas for innovation arrive to random recipients at random times. By foregoing investment in a current idea, society as a whole preserves an option to invest in a better idea for the same market niche, but with delay. Because successive ideas may occur to different people, there is a conflict between private and social optimality. We investigate the optimal policy when the social planner learns over time about the arrival rate of ideas, and when private recipients of ideas can bank their ideas for future use. We argue that private incentives to create socially valuable options can be achieved by giving higher rewards where "ideas are scarce."

1 Introduction

Creativity is an essential ingredient to R&D and to economic growth. When someone conceives of a new gadget to put under the Christmas tree, that is an act of creativity. When someone realizes that light bulbs would waste less energy if light is produced with fluorescence, that is also an act of creativity. These acts of creativity depend on the social environment of the creator, but conditional on the social milieu, have a large random component. Oddly, it is hard to find creativity in the economists' tool kit for studying R&D. In this paper, we suggest a direction of research that might fill that gap.

Instead of focussing on new products, we focus on acts of scientific creativity that suggest ways to satisfy a need. Taking the social milieu as given, we consider innovative environments where agents randomly receive ideas for how to fill a given market niche. Our modeling approach links creativity with scarcity of ideas. In environments where ideas are scarce, only the very creative receive ideas. An agent can invest in his idea or not. Investing provides him with profit, depending on the reward system, but not investing preserves a social option: Another agent with an even better idea might invest instead. This change in perspective expands the purpose of the R&D reward structure. A good reward structure must not only motivate effort and cover the cost of innovation, but must also ensure that ideas are discarded when it is better from society's point of view to preserve the social option.

We study a model where ideas, which represent investment opportunities, are private information and arrive to random agents at random times. This premise diverges from many R&D models in which investment opportunities are common knowledge and eternally present.[1] In those models, if progress is slow, it is because R&D is costly and resources are scarce. In our model, there is another reason that progress can be slow: Ideas for investment are themselves scarce, not only from an individual's point of view, but also for society as a whole. Whether an idea should be used or discarded depends both on the costliness of

---

[1] This is implicitly the premise of a large literature on patent races that builds on models surveyed by Reinganum (1989). These models have also been used in the growth literature. See, for example, Aghion and Howitt (1992).

investing in it and on the rate at which new ideas will arrive. If other ideas will arrive rapidly, it might be best to discard the idea and preserve the option for society to invest in a better idea.

Our objective is to show how rewards should reflect the scarcity of ideas, understood as their arrival rate. The social planner does not know the ideas that have arrived, and does not know who received them. If all ideas were available at the same time, the goal of the social planner would be to find the minimum cost idea. However, that is not possible because the ideas arrive at random times. The planner can weed out high-cost ideas by offering limited rewards, but he still faces a trade-off between cost and delay. To ensure that the market niche is filled at low cost, he may have to endure a costly delay.

Our conclusions about the optimal reward structure are tied to the notion that ideas are scarce, so that the reward policy must mediate between cost and delay. If the social planner knows the scarcity (arrival rate) of ideas, optimal rewards should increase with the scarcity of ideas. Society should be willing to tolerate higher cost in environments where ideas are scarce (the arrival rate is low). However, this policy prescription relies on the strong assumption that the social planner knows the scarcity of ideas. How should the social planner set the optimal policy if the scarcity of ideas is not known? We show that optimal rewards should increase with delay in filling the market niche. Since the social planner is not the recipient of ideas, he must make inferences about the arrival rate of ideas as time passes. As time passes with no innovation, the planner becomes more pessimistic about the arrival rate, and will tolerate higher cost in order to reduce delay. We do not know of other papers where the reward policy changes dynamically.[2]

Our model also implies that the profit on R&D investments will be positive in equilibrium due to the scarcity of ideas. Since investment opportunities are not common knowledge,

---

[2] A dynamically changing reward policy would presumably be optimal in any model where learning takes place about something relevant to rewards. Although the authors do not analyze this aspect, two papers where that might be true are Choi (1991) and Malueg and Tsutsui (1997). In those papers, there is an unknown parameter that governs the hazard rate of success in a production function for R&D that is common knowledge among the firms in a race. In our model below, the planner is learning about the hazard rate at which the population as a whole receives ideas for investment. There is no commonly known but uncertain production function for R&D.

profits are not dissipated.

As in O'Donoghue, Scotchmer and Thisse (OST, 1998) and Scotchmer (1999), our model distinguishes between ideas and innovations. However, delay is never optimal in these papers. OST (1998) address environments where the ideas are complements in the sense that each idea builds on previous ideas, and Scotchmer (1999) addresses environments where ideas serve different market niches, but there are no substitute ideas for a given market niche. In the model we discuss here, it is because ideas are substitutes that a certain amount of delay should be tolerated. One of the ideas that arrives during the delay may have low cost.

Our model is a real options model in the spirit of MacDonald and Siegel (1986) and Dixit and Pindyck (1994). An investment is irreversible and could turn out to be a mistake. To avoid mistakes, there is a value to delay. In many real options models, the value of the option is internalized by the firm. In our model, ideas (investment opportunities) accrue to random firms, which means that although waiting is valuable to society, the value of waiting is not internalized by any potential innovator. The problem of the social planner is to ensure that private recipients of ideas preserve socially optimal options.

Our modeling apparatus is reminiscent of search models (see McCall and McCall, 2008), although we do not interpret our random process as search, and there is an important difference between our social planner and a job searcher. In search models, all opportunities arrive to a single searcher who sets an optimal stopping policy. In our model, ideas are so scarce that no individual is likely to receive more than one idea. The planner who sets the reward policy does not know who receives ideas and when. Despite this fundamental difference, if the social planner knows the arrival rate of ideas, the optimal reward policy is similar to the stopping rule that emerges in search models. However, our main focus is the more realistic case where the arrival rate of ideas is not known. Although some search models involve learning,[3] the case where the searcher is learning about the arrival rate of offers has not been studied in the search literature. Moreover, the inference problem of a job searcher in this case would be fundamentally different. The planner who sets the reward

---

[3] See, for example, Rothschild (1974) and Rosenfield and Shapiro (1981), where there is learning about the distribution of prices.

policy does not know the arrival rate of ideas, does not know who received ideas, and does not know how many ideas have been received. To choose an optimal policy, he must infer something about the arrival rate of ideas using the only information available to him, which is the time that has passed without success. In contrast, a job searcher would know how many job offers have been received and rejected.

The paper is structured as follows. In section 2, we set forth a simple model of scarce ideas. In sections 3 and 4, we characterize the optimal cost threshold that the planner would like to implement at each date. The optimal cost threshold is a cost such that the possessor of a lower-cost idea should invest. The planner realizes that, if someone invests, society is giving up an option, namely, the option to wait for a better idea.

The stochastic process that determines the option value depends on whether rejected ideas are lost forever ("use it or lose it") or banked for future use ("use it or bank it"). We treat these two cases in sections 3 and 4. Banking is attractive to the social planner when he does not know the arrival rate of ideas. As time passes, the planner becomes more pessimistic about the arrival rate. An idea that seemed too costly a year ago will seem more attractive at present because more delay is predicted. The planner will therefore want access to the banked idea with lowest cost.

In section 5, we show how the planner can implement the optimal cost threshold with and without banking of ideas. The optimal reward policy plays on the fact that banking is attractive to the recipient of an idea whenever the reward function is increasing. The recipient of the idea may be willing to forego the profit available by investing at present in order to gamble on a higher reward in the future. Of course, the recipient may be preempted in the meantime.

We conclude in section 6 by mentioning some ways that the optimal reward policy corresponds to legal institutions.

2 Ideas and Innovations

We assume there is a market niche that may be filled with an innovation. The social value of filling the market niche is v/r, where r is the discount rate. There is an exogenous process by which the potential innovators receive ideas for filling the market niche. To innovate, the inventor must first have an idea, which we interpret as an act of creativity, and then have an incentive to invest in it.

Each idea occurs at a random time, to a random recipient. Each idea has associated to it an R&D cost that is drawn independently from a common distribution F with support in $[0, \infty)$ and density f. To create an innovation, the recipient of an idea must invest the cost. We assume that the ideas rain down on the population as a whole according to a Poisson process with parameter λ, and we take the parameter λ as a measure of scarcity. If the hit rate λ is low, ideas are scarce.

We do not assume that there is a cost to creating the flow of ideas. Such an assumption would not change our results in any interesting way, but that is not why we reject it. We reject it because it obscures an important point that we wish to emphasize.

In order to study R&D, or any other incentive problem, one must begin from a model of what is primitive in the economy and what is endogenous. By logical necessity, something is primitive, that is, not chosen. Since the point of R&D is to create or discover knowledge, the thing that is primitive must be the possibilities for how to do so, for example, some notion of what to invest in, and what the investment might lead to. This is generally called a "production function for knowledge," and it is often assumed to be common knowledge. No one explains where the production function comes from; it is simply a primitive of the economy, and is eternally present. Here we assume that the primitive is a process of receiving ideas for R&D investments, and that the ideas are private.[4] This is meant to capture a

---

[4]Berliant and Fujita (2008) make a different departure from the production-function model. They assume that all knowledge is eternally present (for example, in the laws of physics), but locked up in boxes. Each box contains a label (an "idea") describing what can be unlocked if the agent invests a cost. The boxes and labels (investment opportunities) are common knowledge, but the knowledge in the box can only be accessed at a cost, and it remains private unless shared. The investment problem is how to collaborate in opening boxes. The relevant resource constraint is the researcher's time.

rudimentary notion of what it means to have scientific creativity. An idea is represented by a cost, and is not eternally present. It occurs to a random person at a random time. In a more complicated model, the idea could be represented by a production function, or the nature of the process could depend on previous accumulated knowledge (as in OST, 1998). Whether the idea is represented by a cost or a production function, our emphasis is on the timing. If ideas occur rapidly, and if they are substitutes in the sense that they are all for the same market niche, then it makes sense to invest in only one of them and it is intuitive that society should be more discriminating about which ideas to invest in.

If we assumed that there is a cost to turning up ideas (in addition to the cost of turning the ideas into innovations), there would have to be a decision on whether to invest the cost, and then the ideas process would no longer be primitive. The primitive would be the meta-idea of investing in a costly ideas process. The meta-idea would be just another, slightly complicated production function for knowledge whose provenance is not explained. We cut through all this by taking the ideas themselves as primitive. We wish to explore the hypothesis that there is such a thing as "creativity" (having ideas), and that creativity is rewarded in the economy. As we will see, there is a reward to the scarcity of ideas.

We now come to the particulars of the model. We assume that each agent receives at most one idea. This is an intentionally extreme assumption that highlights the main premise of the paper. Ideas are scarce, not only for society as a whole, but especially from the perspective of any individual. The recipient of an idea can invest in it, discard it, or bank it, which means to remember it for future use. If the recipient of an idea invests in it, the process stops because the market niche has been filled. The optimal policy will therefore operate by getting the population of potential innovators to screen their ideas and then to discard or bank those with costs that are too high.

The social option preserved by not investing has value because another idea might be less costly. There is thus a social trade-off between cost and delay. However, this social trade-off is not the private trade-off, because the next idea will occur to someone else. Because an individual is not likely to receive another idea, the individual might be too willing to invest

in the idea he possesses. He is concerned with his private profit, not with the social value of preserving options. The policy challenge is to manage private incentives to invest in a way that is socially optimal.

The social policy is described by a *threshold function* $c : \mathbf{R}_+ \rightarrow \mathbf{R}_+$ such that the recipient of an idea at time t invests if the cost of the idea is less than $c(t)$. We say an idea at time t is *viable* if it has cost less than $c(t)$. The expected cost of a random viable idea that arrives at time t is

$$E_F(c(t)) = \int_0^{c(t)} \hat{c} \frac{f(\hat{c})}{F(c(t))} d\hat{c}. \tag{1}$$

The threshold function c is a stopping rule. We say that the investment process *survives* to t if there is no viable idea before t.

A threshold function c cannot be implemented directly by a social planner, because the social planner does not observe the ideas that are received or who receives them. Much of our task below will consist in showing how the social planner can implement the threshold function that is best for society as a whole. First, however, we must describe the optimal c.

We consider two versions of the ideas process. In the "use it or lose it" model, an idea that is not used immediately is lost. For example, an idea may be lost or forgotten if the recipient moves on to other projects. However, not all ideas will be lost, especially if there is an incentive to remember them. We consider this in the "use it or bank it" model.

3 Use it or lose it

In this section, we assume that, if the recipient of an idea decides not to invest, the idea is lost to everyone, including the recipient, and cannot be reclaimed later.

Let $P(t|\lambda, c)$ be the probability of surviving to time t, as seen from time 0, when the threshold function is c and the arrival rate of ideas is λ. The survival probability differs according to whether ideas can be banked, but in both models, the probability distribution on survival times is stochastically larger at smaller arrival rates.

When recipients either use their ideas or forget them immediately, the instantaneous arrival rate of viable ideas at time t is $\lambda F(c(t))$. As seen from time $t = 0$, the probability

7

of survival to time t is $P(t|\lambda, c)$, defined as

$$P(t|\lambda, c) = e^{-\Lambda(\lambda, t, c)} \quad \text{where} \quad \Lambda(\lambda, t, c) = \int_0^t \lambda F\left(c\left(\tilde{t}\right)\right) d\tilde{t} \tag{2}$$

(See, for example, Snyder and Miller, 1991, p. 51.) The probability of surviving to $\hat{t}$, conditional on surviving to an earlier time t, is $P\left(\hat{t}|\lambda, c\right)/P(t|\lambda, c)$. As seen from time t, and conditional on the arrival rate λ, the probability that the first viable idea arrives at $\hat{t} > t$ is

$$\frac{d}{d\hat{t}}\left[1 - \frac{P\left(\hat{t}|\lambda, c\right)}{P(t|\lambda, c)}\right] = \lambda F\left(c\left(\hat{t}\right)\right) e^{-\left[\Lambda(\lambda, \hat{t}, c) - \Lambda(\lambda, t, c)\right]} \tag{3}$$

Although our focus is on the problem that λ is unknown, we first show how results from the job-search literature, where λ is known, can be applied to our R&D problem. Conditional on λ, and on an arbitrary threshold function c, social welfare measured from time t is V, defined by

$$\begin{aligned} V(t, c, \lambda) &= \int_t^\infty e^{-r(\hat{t} - t)}\left(\frac{v}{r} - E_F\left(c\left(\hat{t}\right)\right)\right) \lambda F\left(c\left(\hat{t}\right)\right) e^{-\left[\Lambda(\lambda, \hat{t}, c) - \Lambda(\lambda, t, c)\right]} d\hat{t} \\ &= \int_t^\infty e^{-r(\hat{t} - t)}\left(\frac{v}{r} - E_F\left(c\left(\hat{t}\right)\right)\right) \lambda F\left(c\left(\hat{t}\right)\right) \frac{P\left(\hat{t}|\lambda, c\right)}{P(t|\lambda, c)} d\hat{t} \end{aligned} \tag{4}$$

This expression reveals the trade-off faced by the policy maker. A higher cost threshold choice at t would imply less delay, since the hit rate of viable ideas, $\lambda F(c(t))$, is higher, but also a higher expected cost of innovation.

If the threshold function c is optimal, the following condition holds at each t.

$$\frac{v}{r} - c(t) = V(t, c, \lambda) \tag{5}$$

The left hand side is the net social value of investing in a threshold idea at time t. The right hand side is the expected, discounted value of waiting for a better idea. By analogy with the job-search literature, the optimal threshold is a stationary value $\bar{c}_\lambda \in \mathbf{R}_+$, such that $c(t) = \bar{c}_\lambda$ for all t, that satisfies

$$\frac{v}{r} - \bar{c}_\lambda = \left(\frac{v}{r} - E_F\left(\bar{c}_\lambda\right)\right) \frac{\lambda F\left(\bar{c}_\lambda\right)}{\lambda F\left(\bar{c}_\lambda\right) + r} \tag{6}$$

A threshold idea with cost $\bar{c}_\lambda$ must generate as much social value (the lefthand side) as waiting for another viable idea (the righthand side). The next idea that meets the cost threshold will be cheaper in expectation ($E_F(\bar{c}_\lambda)$ instead of $\bar{c}_\lambda$), but that benefit must be discounted $\left(\frac{\lambda F(\bar{c})}{(\lambda F(\bar{c})+r)} < 1\right)$.

The optimality condition (6) has an interesting and intuitive interpretation, taking λ as a measure of scarcity. If ideas are scarce in the sense that they occur at long intervals (λ is low), then society should be willing to bear a higher cost to fill the market niche than when ideas occur frequently. When λ is high, the cost of waiting is reduced, and it is optimal to be more selective in choosing an idea for investment. This is stated in the following remark.[5]

Remark 1 *Suppose that the arrival rate of ideas, λ, is fixed and known. Then the optimal cost threshold is a stationary value that depends on λ, and the optimal stationary threshold is decreasing with λ.*

We now turn to the more realistic case that λ is unknown. Like all contracts, R&D incentives must depend on things that are verifiable. A prize or patent authority knows whether the market niche has been filled, but does not observe the hypothetical distribution of arrival times, and does not observe the arrival of ideas that are rejected. The length of time without arrival of a viable idea is a signal of λ. A long period with no arrival should make the social planner more pessimistic about λ – it shifts the posterior distribution on λ toward lower values. However, the posterior distribution on λ must also account for the fact that some ideas are rejected. Thus, the threshold function for accepting or rejecting ideas is an ingredient to forming a posterior belief on λ.[6]

We show that, when the posterior distribution on λ is changing as time passes, neither the value function (optimized welfare function) nor the optimal investment strategy

---

[5] We state the result as a remark because it follows from the job-search literature. We prove it in Erkal and Scotchmer (2009).

[6] This is where the R&D problem differs importantly from the job-search problem. A job searcher can update his belief on λ by using the number of offers that are received. Because he knows of the job offers directly, the threshold wage function by which he would have accepted or rejected offers in the past is irrelevant.

is stationary. Because the posterior distribution on λ shifts toward lower values as time passes with no viable idea, the (optimized) value of waiting for a better idea decreases with time. This implies that society should optimally be less discriminating about which idea is accepted. In particular, the socially optimal cost threshold is increasing instead of being stationary.

Let $\tilde{h}$ be the prior density function for the distribution of λ with support $[0, \infty)$. Then the posterior density, conditional on a threshold function c, and conditional on no viable hit having arrived by time t, is $h\left(\cdot|t, c\right)$ with cumulative distribution $H\left(\cdot|t, c\right)$, where $h\left(\cdot|t, c\right)$ satisfies

$$h\left(\lambda|t, c\right) = \frac{\tilde{h}\left(\lambda\right) e^{-\Lambda(t, c, \lambda)}}{\int \tilde{h}\left(\lambda\right) e^{-\Lambda(t, c, \lambda)} d\lambda} \quad \text{for each } \lambda \in (0, \infty) \tag{7}$$

Let $E\left(\lambda|t, c\right)$ be the expected value of λ:

$$E\left(\lambda|t, c\right) = \int_0^\infty \lambda h\left(\lambda|t, c\right) d\lambda$$

In the appendix we prove the following lemma:

Lemma 1 [Increasing pessimism] *If $t_1 < t_2$, the distribution $H\left(\cdot|t_1, c\right)$ stochastically dominates $H\left(\cdot|t_2, c\right)$. Moreover, $E\left(\lambda|t, c\right)$ decreases with t.*

Let $\phi\left(\hat{t}, t, c\right)$ be the expected probability of stopping (investing) at $\hat{t}$, as seen from time t, when the stopping rule (threshold function) is c and the beliefs are $h\left(\cdot|t, c\right)$:

$$\phi\left(\hat{t}, t, c\right) \equiv \int_0^\infty \lambda F\left(c\left(\hat{t}\right)\right) e^{-\left[\Lambda\left(\hat{t}, c, \lambda\right) - \Lambda(t, c, \lambda)\right]} h\left(\lambda|t, c\right) \, d\lambda$$

The social value of continuing from time t is given by a function $\tilde{V}$, defined in the first line of (8). Substituting for V from (4) and using the definition of ϕ gives us the second line.

$$\tilde{V}\left(t, c, \tilde{h}\right) = \int_0^\infty V\left(t, c, \lambda\right) h\left(\lambda|t, c\right) \, d\lambda \tag{8}$$
$$= \int_t^\infty e^{-r\left(\hat{t}-t\right)} \left(\frac{v}{r} - E_F\left(c\left(\hat{t}\right)\right)\right) \phi\left(\hat{t}, t, c\right) d\hat{t}$$

Our objective is to show that $\tilde{V}\left(t, c, \tilde{h}\right)$ is decreasing with t when the threshold function c is optimally chosen. We need the following derivatives, which show that, starting from $t + dt$ instead of t, the probability of stopping at each later date $\hat{t}$ is larger.

10

$$\frac{d}{dt} h\left(\lambda | c, t\right) = h\left(\lambda | c, t\right) F\left(c\left(t\right)\right) \left[E\left(\lambda | t, c\right) - \lambda\right]$$

$$
\begin{aligned}
\frac{d}{dt} \phi\left(\hat{t}, t, c\right) &= \int_0^\infty \lambda^2 F\left(c\left(\hat{t}\right)\right) F\left(c\left(t\right)\right) e^{-\left[\Lambda\left(\hat{t}, c, \lambda\right) - \Lambda\left(t, c, \lambda\right)\right]} h\left(\lambda | t, c\right) d\lambda \\
&\quad + \int_0^\infty \lambda F\left(c\left(\hat{t}\right)\right) e^{-\left[\Lambda\left(t, c, \lambda\right) - \Lambda\left(t, c, \lambda\right)\right]} \frac{d}{dt} h\left(\lambda | t, c\right) d\lambda \\
&= E\left(\lambda | t, c\right) F\left(c\left(t\right)\right) \phi\left(\hat{t}, t, c\right)
\end{aligned}
$$

Then[7]

$$
\begin{aligned}
&\frac{d}{dt} \tilde{V}\left(t, c, \tilde{h}\right) \\
&= r\tilde{V}\left(t, c, \tilde{h}\right) + E\left(\lambda | t, c\right) F\left(c\left(t\right)\right) \left(\tilde{V}\left(t, c, \tilde{h}\right) - \left(\frac{v}{r} - E_F\left(c\left(t\right)\right)\right)\right) \quad (9) \\
&= \left[r + E\left(\lambda | t, c\right) F\left(c\left(t\right)\right)\right] \tilde{V}\left(t, c, \tilde{h}\right) - E\left(\lambda | t, c\right) F\left(c\left(t\right)\right) \left(\frac{v}{r} - E_F\left(c\left(t\right)\right)\right) \quad (10)
\end{aligned}
$$

Let $c : \mathbf{R}_+ \to \mathbf{R}_+$ be the threshold function that maximizes $\tilde{V}\left(0, \cdot, \tilde{h}\right)$. At a given time t, the planner's choice is the function c on the domain $[t, \infty)$. However, the entire function c is relevant for $\tilde{V}$ at t. The values of the function c on the domain $[0, t)$ affect the value $\tilde{V}\left(t, c, \tilde{h}\right)$ because they determine the probability that ideas were received and rejected, and thus determine $h\left(\cdot | t, c\right)$.

The optimal threshold function c requires that the stopping value is equal to the continuation value:

$$\frac{v}{r} - c\left(t\right) = \tilde{V}\left(t, c, \tilde{h}\right), \qquad \text{each } t \in (0, \infty) \qquad (11)$$

Proposition 1 *Suppose that the recipient of an idea must use it or lose it. Suppose that the arrival rate of ideas, λ, has a prior distribution $\tilde{h}$ with support $[0, \infty)$. Let c be the threshold function that maximizes $\tilde{V}\left(0, \cdot, \tilde{h}\right)$. Then c is increasing.*

---

[7]In dynamic optimization problems, the optimality condition can be expressed as if the continuation value is a capital asset that deserves a normal rate of return. The derivative can be derived using this principle. The return is

$$r\tilde{V} dt = E\left(\lambda | t, c\right) F\left(c\left(t\right)\right) \left(\frac{v}{r} - E_F\left(c\left(\hat{t}\right)\right) - \tilde{V}\right) dt + \frac{d}{dt} \tilde{V} dt$$

Intuitively, V is decreasing because the observer becomes more and more pessimistic about the arrival rate of ideas as time continues without a viable hit. Because of this pessimism, more delay is expected. To mitigate delay, it is optimal to tolerate higher cost.

In section 5, we discuss how the optimal cost threshold can be implemented. In the "use it or lose it" model, it is easy to implement the optimal cost threshold by setting the reward equal to the cost threshold. This is because each recipient of an idea has a single opportunity to invest. He will not receive another idea (ideas are scarce), and he must either invest in the idea immediately or lose it forever.

4   Use it or Bank it

When the reward is equal to the cost threshold, and therefore (with unknown λ) increasing, the possessor of an idea may have an incentive to delay investment to get a higher reward. If the recipient can bank his idea for later use, the social planner needs to take this into account in choosing the optimal cost threshold as well as a reward policy to implement it.

How should the planner view banking? The social planner does not want to delay investments that should be viable under his optimal cost threshold. His reward function should ensure that this does not happen. At the same time, banking ideas for future use is tantamount to increasing the arrival rate of ideas in the future. Since this is valuable, the optimal cost threshold should take it into account.

With banking, the social policy is again a threshold function c. Ideas accumulate over time and are banked by the recipients. An idea that is converted to an innovation can either be a banked idea or a new idea. As before, we first consider the case that λ is known, and then consider the optimal threshold function when the beliefs about λ are evolving.

The probability of stopping (investing) at time t must be described differently according to whether the threshold function c is increasing or decreasing at that t. If decreasing, the banked ideas are irrelevant. Any banked idea that would be chosen at t would also have been chosen at $t - dt$. If there is investment at t, it is because a viable idea materializes at that moment. On the other hand, if c is increasing at t, then banked ideas may become

Figure 1: The stochastic process with banking

viable. If c is increasing, both the banked ideas and the increasing cost threshold affect the probability of investment at time t.

We will describe the probabilities of survival at each t by reference to Figure 1, which shows an arbitrary threshold function. To describe the probability of investing in an idea at any t where c is decreasing, such as t_1 in Figure 1, let t_m be the largest value smaller than t where c is nonincreasing. During the time period $[t_m, t]$, no banked ideas become viable. Let $Q(t_m|\lambda, c)$ be the probability of surviving to t_m. Then the probability of survival to t, (that is, the probability that there is no viable idea by time t) is similar to the case without banking:

$$Q(t_m|\lambda, c) e^{-\Lambda(\lambda, t, c) + \Lambda(\lambda, t_m, c)}$$

This is the probability of survival to t_m times the probability that no viable idea arrives in the interval $[t_m, t]$ where c is decreasing.

To describe the probabilities of survival at t where c is increasing, let $t_0(t)$ be the largest

value smaller than t such that $c(t_0(t)) = c(t)$. If there is no such value, let $t_0(t) = 0$.

In Figure 1, c is increasing at t_2 and t_3. At t_2, the relevant banked ideas have been accumulating for a shorter period of time than at t_3. At t_2, any ideas below the cost threshold $c(t_2)$ that were received before $t_0(t_2)$ would have been used before $t_0(t_2)$. Therefore the banked ideas that might become viable are those which accumulated between $t_0(t_2)$ and t_2. At t_3, there may be relevant ideas with cost near $c(t_3)$ that accumulated very early, since there was never a time when such high-cost ideas were below the cost threshold.

The probability of survival to t (that is, the probability that there is no viable idea by time t) is

$$Q(t_0(t)|\lambda, c)e^{-\lambda F(c(t))[t - t_0(t)]}$$

This is the probability of survival to $t_0(t)$ times the probability that no viable idea arrives in the interval $[t_0(t), t]$, where c is increasing at t.

Thus, if ideas are banked, the stochastic process that determines the probability of survival until t satisfies

$$Q(t|\lambda, c) = \begin{cases} Q(t_m|\lambda, c) e^{-[\Lambda(\lambda, t, c) - \Lambda(\lambda, t_m, c)]} & \text{if} \quad c \text{ is decreasing in } [t_m, t] \\ Q(t_0(t)|\lambda, c)e^{-\lambda F(c(t))[t - t_0(t)]} & \text{if} \quad c \text{ is increasing at } t \end{cases} \tag{12}$$

As seen from time t, the probability of arriving at $\hat{t}$ is $\frac{Q(\hat{t}|\lambda, c)}{Q(t|\lambda, c)}$. The probability that the first viable idea becomes available at $\hat{t}$ is the probability of arriving there, times the instantaneous probability that a viable idea arrives at time $\hat{t}$, namely,

$$\frac{d}{d\hat{t}}\left[1 - \frac{Q(\hat{t}|\lambda, c)}{Q(t|\lambda, c)}\right] = \lambda \mathcal{F}(\hat{t}|c) \frac{Q(\hat{t}|\lambda, c)}{Q(t|\lambda, c)}$$

where

$$\mathcal{F}(t|c) = \begin{cases} F(c(t)) & \text{if } c'(t) \leq 0 \\ [F(c(t)) + f(c(t))(t - t_0(t))c'(t)] & \text{if } c'(t) > 0 \end{cases}$$

When c is decreasing, the instantaneous probability of an innovation, $\lambda \mathcal{F}(\hat{t}|c) = \lambda F(c(\hat{t}))$, is the same as in the "use it or lose it" model, namely, the probability that a viable idea occurs in the interval $(\hat{t}, \hat{t} + dt)$. But when c is increasing, the innovation may result from

14

a banked idea rather than from an idea that occurs in the interval $(\hat{t}, \hat{t} + dt)$. The instantaneous probability of innovation is therefore larger, namely, $\lambda \mathcal{F}(\hat{t}|c) > \lambda F(c(\hat{t}))$. The instantaneous probability of innovation has two parts. First is the probability that a viable idea arrives to someone in dt, namely $\lambda F(c(\hat{t})) dt$. Second is the probability that a banked idea becomes viable. When the threshold rises by $c'(\hat{t}) dt$, the probability that there is a banked idea in the cost band $c'(\hat{t}) dt$ is $\lambda f(c(\hat{t})) (\hat{t} - t_0(\hat{t}))$.

Conditional on surviving to time t, and conditional on λ, social welfare measured from time t is B defined by

$$B(t, c, \lambda) = \int_t^\infty e^{-r(\hat{t} - t)} \left(\frac{v}{r} - E_F(c(\hat{t})) \right) \lambda \mathcal{F}(\hat{t}|c) \frac{Q(\hat{t}|\lambda, c)}{Q(t|\lambda, c)} d\hat{t}$$

If the threshold function c is optimal and λ known, the following condition holds at each t.

$$\left(\frac{v}{r} - c(t) \right) = B(t, c, \lambda) \tag{13}$$

When ideas can be banked and λ is known, it remains true that the optimal cost threshold is stationary, as in the job-search model. If the threshold is stationary, then the probability distribution $Q(\cdot|\lambda, \bar{c})$ is the same as $P(\cdot|\lambda, \bar{c})$, so $\bar{B}(t, \bar{c}, \lambda) = \bar{V}(t, \bar{c}, \lambda)$ and the optimal cost threshold is the same in both cases. Thus, Remark 1 holds for both the "use it or lose it" model and the "use it or bank it" model.[8]

However, we are mainly interested in the case that λ is unknown. The prior is again $\tilde{h}$, and using the survival probabilities described in (12), the posterior distribution on λ is again described by (7), substituting Q for P. The analog to Lemma 1 holds for the distribution Q, by the same proof as for the distribution P.[9] $E(\lambda|\hat{t}, c)$ decreases with $\hat{t}$.

Let $\beta(\hat{t}, t, c)$ be the expected probability of stopping at $\hat{t}$, as seen from time t, when the

---

[8]Banking is called "recall" in the job-search literature. McCall and McCall (2008) give a very high-level proof of stationarity for the case of recall, relying on the fact that the first-order condition has a unique solution. In Erkal and Scotchmer (2009), we give a more direct proof using our social welfare function.

[9]Lemma 1 is proved by using Claim 1 in the proof. Claim 1 also applies here except that, for the distribution Q,

$$\frac{d}{dt} h(\lambda|t, c) = \mathcal{F}(c(t)) h(\lambda|t, c) [E(\lambda|t, c) - \lambda].$$

threshold function is c and the beliefs are $h\left(\cdot|t,c\right)$:

$$\beta\left(\hat{t},t,c\right)\equiv\int_{0}^{\infty}\lambda\mathcal{F}\left(\hat{t}|c\right)\frac{Q\left(\hat{t}|\lambda,c\right)}{Q\left(t|\lambda,c\right)}h\left(\lambda|t,c\right)\ d\lambda$$

Differentiating,

$$\frac{d}{dt}\beta\left(\hat{t},t,c\right)=E\left(\lambda|t,c\right)\mathcal{F}\left(t|c\right)\beta\left(\hat{t},t,c\right)$$

The social value of continuing from time t is given by a function $\tilde{B}$, which we write in two ways. The first line is the definition, and the second line is equivalent, emphasizing that the belief on λ is updated at each $\hat{t}$.

$$\begin{aligned}
\tilde{B}\left(t,c,\tilde{h}\right) &= \int_{0}^{\infty}B\left(t,c,\lambda\right)h\left(\lambda|t,c\right)d\lambda & (14)\\
&= \int_{t}^{\infty}e^{-r\left(\hat{t}-t\right)}\left(\frac{v}{r}-E_{F}\left(c\left(\hat{t}\right)\right)\right)\beta\left(\hat{t},t,c\right)d\hat{t}
\end{aligned}$$

Then

$$\begin{aligned}
\frac{d}{dt}\tilde{B}\left(t,c,\tilde{h}\right) &=\\
&= r\tilde{B}\left(t,c,\tilde{h}\right)+E\left(\lambda|t,c\right)\mathcal{F}\left(t|c\right)\left[\tilde{B}\left(t,c,\tilde{h}\right)-\left(\frac{v}{r}-E_{F}\left(c\left(t\right)\right)\right)\right] & (15)\\
&= \left[r+E\left(\lambda|t,c\right)\mathcal{F}\left(t|c\right)\right]\tilde{B}\left(t,c,\tilde{h}\right)-E\left(\lambda|t,c\right)\mathcal{F}\left(t|c\right)\left(\frac{v}{r}-E_{F}\left(c\left(t\right)\right)\right) & (16)
\end{aligned}$$

As in the model without banking, the entire function c is relevant for $\tilde{B}$ at t. The values of the function c on the domain $[0,t)$ affect the value $\tilde{B}\left(t,c,\tilde{h}\right)$ because they determine the probability that ideas were received and rejected, and thus enter the beliefs $h\left(\cdot|t,c\right)$.

The optimal c satisfies

$$\frac{v}{r}-c\left(t\right)=\tilde{B}\left(t,c,\tilde{h}\right) \tag{17}$$

The following proposition, proved in the appendix, shows that the optimal cost threshold with banking is again increasing.

Proposition 2 *Suppose that the recipient of an idea can use it or bank it. Suppose that the arrival rate of ideas, λ, has a prior distribution $\tilde{h}$ with support $[0,\infty)$. Let c be the threshold function that maximizes $\tilde{B}\left(0,\cdot,\tilde{h}\right)$. Then c is increasing.*

16

Finally, we show that the optimized social welfare is higher with banking than without. When ideas are banked, the social planner is more pessimistic about λ at each t for a given c. At the same time, the arrival rate of viable ideas is higher when some ideas may come from the idea bank. The next proposition shows that the latter effect dominates.

Proposition 3 *Let $\mathcal{C}$ be the set of threshold functions c that are increasing. Then for each $t > 0$, $\max_{c \in \mathcal{C}} \tilde{B}\left(t, c, \tilde{h}\right) > \max_{c \in \mathcal{C}} \tilde{V}\left(t, c, \tilde{h}\right)$.*

From (11) and (17), this proposition implies the following:

Corollary 1 *Let c^V be the optimal threshold function in the "use it or lose it" model, and let c^B be the optimal threshold function in the "use it or bank it" model. Then $c^V > c^B$.*

The social planner prefers to be more selective when he can rely on banked ideas, even if he is more pessimistic about λ at each t.

5   Implementing the Optimal Cost Threshold

The social planner cannot implement the optimal cost threshold directly, because the social planner is not the recipient of the ideas. Ideas for R&D are widely dispersed within the population of potential innovators. At best the social planner can try to implement the optimal threshold by setting rewards.

We suppose that the social planner sets a *reward function* $\rho : \mathbf{R}_+ \rightarrow \mathbf{R}_+$. For example, the reward function can represent patent policy or a prize system. The reward function ρ *implements* the threshold function c if the possessor of an idea with cost c_0 at time t invests in the idea if and only if his idea satisfies $c_0 \leq c(t)$. In the "use it or lose it" model, the only relevant ideas are those that just arrived, but in the "lose it or bank it" model, the relevant idea might previously have been banked.

It is trivial to implement the optimal cost threshold (or any other cost threshold) in the "use it or lose it" model, simply by setting $\rho(t) = c(t)$ for all t. Since ideas cannot be banked, the only choice is to invest immediately in return for the reward or to give up

the idea forever. When λ is known, the reward $\rho(t) = c(t)$ works in the "use it or bank it" model as well, because the optimal cost threshold is stationary. It is better to get a stationary reward immediately than with delay. These observations lead to an interesting and intuitive economic interpretation about rewards, recorded in the following proposition: When ideas are scarcer, rewards should be higher. This is immediate in part (a), since scarcity is measured by λ. In part (b), c and ρ are increasing because the social planner becomes more pessimistic about λ as time passes without success. As time passes, ideas seem scarcer, and it is optimal to let the reward rise.

Proposition 4 (Rewards should be higher when ideas are scarcer.)
(a) In either the "use it or lose it" model or the "use it or bank it" model, if λ is fixed and known, the optimal reward is stationary, and is equal to the optimal cost threshold c_λ. The optimal stationary reward is higher when λ is lower.
(b) In the "use it or lose it" model, when λ is unknown, the optimal reward $\rho(t)$ is equal to the optimal cost threshold $c(t)$ at each t, and is increasing.

Implementation is not as easy in the "use it or bank it" model when the arrival rate of ideas is unknown, although, as we will see, rewards should still increase with time.

In both the "use it or lose it" model and the "use it or bank it" model, the optimal cost threshold is increasing with the passage of time when the arrival rate of ideas is unknown. In the "use it or bank it" model, this implies that the reward must be greater than the implemented cost threshold at each t, instead of equal to it. If $\rho(t) = c(t)$ for all t, as in the "use it or lose it" model, the recipient of a threshold idea (cost equal to $c(t)$) would not invest immediately as intended, since investing would lead to zero profit. The possessor of a threshold idea at t might make positive profit by letting some time pass so that the reward is higher. Even if the possessor of the idea might be preempted during the delay, the expected profit with delay is still larger than zero.[10]

We suppose that the social planner chooses a reward function ρ, and the recipients

---

[10] A similar type of trade-off exists in Weeds (2002), who considers a model of R&D competition where delay is undermined by the fear of preemption.

of ideas choose investment strategies. Each recipient's investment strategy is a threshold function that indicates whether, when the opportunity arises, the possessor of the idea will invest in the idea or bank it. The planner's objective is to make sure that the privately chosen threshold functions correspond to the threshold function that is optimal.

Each idea recipient's incentive to bank or invest depends on his belief about λ, and also on his belief about the other agents' investment strategies. If there is a large accumulation of banked ideas, the probability of being preempted is high.

Further, the social planner must predict these beliefs. If the social planner does not know the beliefs of the idea recipients, he cannot predict their investment strategies, and thus cannot predict the cost threshold that will be implemented by his reward function.

We solve the problem of beliefs in a familiar way. We require that beliefs are correct in equilibrium. For the threshold function c that will be implemented, a recipient must believe in equilibrium that other agents invest according to c, and the recipient must find it most profitable to invest according to c himself. The point is to find a reward function with this result.

The planner's belief about λ is irrelevant in the following discussion. The planner can implement any nondecreasing threshold he wants, provided he knows the beliefs of the idea recipients. The idea recipients have more information than the planner, and will have different beliefs than the planner. This is because, when a nonviable idea arrives to a recipient, the arrival contains information about λ even if the idea is banked instead of used.

Let $\hat{h}(\cdot|t, c)$ represent the belief of each recipient about λ, with expected value $\hat{E}(\lambda|t, c)$. The argument c is a belief, namely, the recipient's belief about the investment strategy (cost threshold) of the other recipients. In the following Remark, we record that if all recipients believe that the other idea recipients obey a common threshold function c, then they also have common beliefs about λ, which we call $\hat{h}$. We then show that all idea recipients will, in fact, choose the same investment strategy in equilibrium, so that these assumptions are justified.

We write the following as a remark instead of a lemma because the proof in the appendix elaborates the model, assuming that the population of idea recipients is finite instead of infinite, and taking limits. We do this in order to derive limit beliefs as the population becomes large. We show that, in the limit, the beliefs of idea recipients do not depend on when a recipient received his idea. In the limit, the probability of receiving an idea is zero, and the timing of the idea has negligible impact beyond the impact of receiving one. Nevertheless, the limit beliefs are more optimistic than those of the planner because the planner has not observed the arrival of any idea at all.

Remark 2 *Let c be an arbitrary nondecreasing investment strategy (cost threshold). (a) If all recipients of ideas believe that c is the investment strategy of every other recipient, then at a given time t, every recipient of an idea has the same belief $\hat{h}(\cdot|t,c)$ on arrival rates, namely*

$$\hat{h}\left(\hat{\lambda}|t,c\right) = \frac{\hat{\lambda}e^{-\hat{\lambda}F(c(t))t}\tilde{h}\left(\hat{\lambda}\right)}{\int \lambda e^{-\lambda F(c(t))t}\tilde{h}\left(\lambda\right)d\lambda} = \frac{\hat{\lambda}Q\left(t|\hat{\lambda},c\right)\tilde{h}\left(\hat{\lambda}\right)}{\int \lambda Q\left(t|\lambda,c\right)\tilde{h}\left(\lambda\right)d\lambda}, \qquad \lambda \in (0,\infty)$$

(b) The recipients' belief $\hat{h}(\cdot|t,c)$ stochastically dominates the planner's belief $h(\cdot|t,c)$ at every t.

From the perspective of an initial time t_0, a recipient's belief $\mathcal{P}$ on the probability of arriving at $t > t_0$ is given by

$$\mathcal{P}(t|c) = \int_0^\infty Q(t|\lambda,c)\,\hat{h}(\lambda|t_0,c)\,d\lambda$$

where c is the threshold function that other recipients are assumed to obey.[11] Using the fact that $\hat{h}$ can be written as

$$\hat{h}(\lambda|t,c) = \frac{Q(t|\lambda,c)\,\hat{h}(\lambda|t_0,c)}{\int Q(t|\lambda,c)\,\hat{h}(\lambda|t_0,c)\,d\lambda}$$

the probability of arriving at $\hat{t}$, having already arrived at t, is given by

$$\frac{\mathcal{P}(\hat{t}|c)}{\mathcal{P}(t|c)} = \int_0^\infty \frac{Q(\hat{t}|\lambda,c)}{Q(t|\lambda,c)}\hat{h}(\lambda|t,c)\,d\lambda$$

---

[11] We use an arbitrary $t_0 > 0$ instead of 0 as the reference point, because no ideas could have been received at time zero.

To define the idea recipient's profit function, suppose that he possesses an idea with cost c_0 at time t. The innovator's profit is a function of the time $\hat{t}$ at which he intends to invest if the market niche has not been filled. This is given by

$$\pi\left(\hat{t}, c_0 | t, c\right) = \left(\rho\left(\hat{t}\right) - c_0\right) e^{-r(\hat{t}-t)} \frac{\mathcal{P}\left(\hat{t}|c\right)}{\mathcal{P}\left(t|c\right)} \tag{18}$$

The derivative of the profit function is

$$
\begin{aligned}
\frac{d}{d\hat{t}} \pi\left(\hat{t}, c_0 | t, c\right) &= \rho'\left(\hat{t}\right) e^{-r(\hat{t}-t)} \frac{\mathcal{P}\left(\hat{t}|c\right)}{\mathcal{P}\left(t|c\right)} - \left(\rho\left(\hat{t}\right) - c_0\right) e^{-r(\hat{t}-t)} \left[r \frac{\mathcal{P}\left(\hat{t}|c\right)}{\mathcal{P}\left(t|c\right)} + \frac{d}{d\hat{t}} \frac{\mathcal{P}\left(\hat{t}|c\right)}{\mathcal{P}\left(t|c\right)} \right] \\
&= e^{-r(\hat{t}-t)} \frac{\mathcal{P}\left(\hat{t}|c\right)}{\mathcal{P}\left(t|c\right)} \times \left(\rho'\left(\hat{t}\right) - \left(\rho\left(\hat{t}\right) - c_0\right) \left[r + \mathcal{F}\left(\hat{t}|c\right) \hat{E}\left(\lambda|\hat{t}, c\right) \right] \right)
\end{aligned}
$$

where

$$\hat{E}\left(\lambda|\hat{t}, c\right) = \int_0^\infty \lambda \hat{h}\left(\lambda|\hat{t}, c\right) d\lambda$$

The optimal investment decision at $\hat{t}$ is

$$
\begin{aligned}
\text{delay if } \rho'\left(\hat{t}\right) &> \left(r + \hat{E}\left(\lambda|\hat{t}, c\right) \mathcal{F}\left(\hat{t}|c\right) \right) \left[\rho\left(\hat{t}\right) - c_0 \right] \\
\text{invest if } \rho'\left(\hat{t}\right) &\leq \left(r + \hat{E}\left(\lambda|\hat{t}, c\right) \mathcal{F}\left(\hat{t}|c\right) \right) \left[\rho\left(\hat{t}\right) - c_0 \right]
\end{aligned} \tag{19}
$$

In (19), ρ' on the left hand side is the benefit of delay. The right hand side is the cost of delay, namely, the interest cost r on the foregone profit $\rho\left(\hat{t}\right) - c_0$, and the perceived probability $\hat{E}\left(\lambda|\hat{t}, c\right) \mathcal{F}\left(\hat{t}|c\right)$ of being preempted.

The investment strategy derived in (19) is the same for all idea recipients, as we initially assumed to derive the beliefs.

The optimal investment behavior (19) should guide the planner in choosing his reward function. For an arbitrary nondecreasing cost threshold c, let the reward function satisfy

$$\rho'(t) = \left(r + \hat{E}\left(\lambda|t, c\right) \mathcal{F}\left(t|c\right) \right) \left[\rho(t) - c(t) \right] \tag{20}$$

The following lemma says that, with the reward function defined in (20), idea recipients will indeed obey the threshold function c that determines beliefs (represented by $\hat{E}\left(\lambda|t, c\right) \mathcal{F}\left(t|c\right)$).

Lemma 2 *Let c be a nondecreasing cost threshold, and suppose that the belief of each idea recipient is c. Suppose that the reward function ρ solves (20) at every t. Then for each t, a recipient's most profitable investment strategy is to invest if he has an idea with cost $c_0 \le c(t)$, and not otherwise.*

Proof: The condition (19) is clearly optimal at the margin, for choosing whether to invest at t or delay for a length of time dt. We must also show that if (19) holds, a longer delay is also not profitable.

If $c_0 \le c(t) \le c(\tilde{t})$, then

$$\rho'(t) = \left(r + \hat{E}(\lambda|t,c) \mathcal{F}(t|c)\right) [\rho(t) - c(t)] \le \left(r + \hat{E}(\lambda|t,c) \mathcal{F}(t|c)\right) [\rho(t) - c_0]$$

and

$$\rho'(\tilde{t}) = \left(r + \hat{E}(\lambda|\tilde{t},c) \mathcal{F}(\tilde{t}|c)\right) [\rho(\tilde{t}) - c(\tilde{t})] \le \left(r + \hat{E}(\lambda|\tilde{t},c) \mathcal{F}(\tilde{t}|c)\right) [\rho(\tilde{t}) - c_0]$$

The first line (respectively, second line) means that it is more profitable to invest at t rather than $t+dt$ (respectively, $\tilde{t}$ rather than $\tilde{t}+dt$) because the additional profit from delay (the left hand term) is no greater than the cost of delay (the right hand term). The first line (respectively, second line) holds because $c_0 \le c(t)$ (respectively, $c_0 \le c(\tilde{t})$). If the idea was not available at t (or if the possessor of the idea made a mistake by not investing), he will invest at the earliest next time, such as at $\tilde{t}$, since $c_0 \le c(t)$ implies $c_0 \le c(\tilde{t})$ whenever $t < \tilde{t}$. At any time after t, the possessor of an idea with cost less than $c(t)$ prefers to invest rather than bank. $\square$

Thus, if c is nondecreasing and the reward function ρ is chosen to satisfy (20), recipients of ideas will invest according to the investment strategy c. We therefore say that ρ implements c if ρ satisfies (20) and also satisfies $\rho(t) \ge c(t)$.

Proposition 5 *Suppose that c is nondecreasing. There exists a reward function ρ such that ρ implements c, $(v/r) \ge \rho > c$ if c is increasing, and $\rho(t) - c(t) \to 0$ as $t \to \infty$.*

Proof: Let

$$\rho(t) = c(t) + \frac{k - \int_0^t e^{-D(\hat{t})} c'(\hat{t})\, d\hat{t}}{e^{-D(t)}} \tag{21}$$

$$\text{where } D(\hat{t}) = \int_0^{\hat{t}} \left(r + \hat{E}(\lambda|\tilde{t}, c)\, \mathcal{F}(\tilde{t}|c) \right) d\tilde{t}$$

$$k = \lim_{t \to \infty} \int_0^t e^{-D(\hat{t})} c'(\hat{t})\, d\hat{t}.$$

The function ρ defined by (21) is a solution to (20). Since $\int_0^t e^{-D(\hat{t})} c'(\hat{t})\, d\hat{t}$ is nondecreasing with t, the choice of k ensures that $\rho(t) - c(t) \geq 0$ for all t, with strict inequality if c is increasing on some domain. The definition of k can be satisfied because at each t, $\int_0^t e^{-D(\hat{t})} c'(\hat{t})\, d\hat{t} \leq \int_0^t c'(\hat{t})\, d\hat{t} = c(t) - c(0)$. Since c is nondecreasing and bounded by (v/r), there exists $\bar{c} \leq (v/r)$ such that $c(t) \to \bar{c}$.

It holds that $\rho(t) - c(t) \to 0$ because

$$\frac{k - \int_0^t e^{-D(\hat{t})} c'(\hat{t})\, d\hat{t}}{e^{-D(t)}} \to 0$$

by L'Hopital's rule:

$$\frac{e^{-D(t)} c'(t)}{e^{-D(t)} \left(r + \hat{E}(\lambda|t, c)\, \mathcal{F}(t|c) \right)} = \frac{c'(t)}{\left(r + \hat{E}(\lambda|t, c)\, \mathcal{F}(t|c) \right)} \to 0$$

Since ρ is increasing by (20), this also shows that $\rho \leq (v/r)$. $\square$

We close this section with a comment on the profitability of R&D in aggregate. Due to the scarcity of ideas, innovators make positive profit on average. This is because the cost of an implemented idea will generally be lower than the threshold. The recipient of a low-cost idea is in a favored position, and everyone would like to have such an idea, but there is little that one can do to create the investment opportunity. In fact, we have taken the extreme assumption that investment opportunities arrive entirely by chance. We have done this to emphasize our key departure from the more standard R&D literature, where all firms have access to an investment opportunity, and profit may be dissipated in a patent race or through preemptive strategies. As a consequence, one would not expect to observe in equilibrium that the return to R&D investments is the same as the return to capital. On average, it should be higher.

Our argument that the optimal cost threshold can be implemented seems to overlook the possibility that an innovator might be able to keep his innovation secret while charging a proprietary price. This would presumably subvert the objective of the reward, and possibly not implement the optimal cost threshold. However, as we show in Erkal and Scotchmer (2007), at least in the case of known λ, secrecy is never optimal.[12] This assumes that with secrecy, another innovator can claim the reward and end the prior innovator's proprietary profit stream.

6 Economic Concepts and Legal Concepts

We interpret ideas, and the fact that ideas are private, as a model of creativity or imagination. Ideas have economic value because they are scarce. We have argued that rewards should be higher in environments where ideas are scarce. If ideas are scarce, higher cost should be tolerated in order to reduce delay. We have also argued that rewards should be increasing as time passes without an innovation. Longer delay leads to expectation of an even longer delay. The delay can be mitigated with higher rewards, since higher rewards encourage investment in higher-cost ideas. Because ideas are not common knowledge, innovators make positive profit in expectation.

These arguments apply equally well to patents and prizes, and any other way of giving rewards.[13] Patents raise the issue of whether our prescriptions can be implemented under existing patent doctrine. They also raise the question of how deadweight loss incurred in collecting the reward money changes the optimal cost threshold.

The main requirements for obtaining a patent are novelty, nonobviousness, utility and enablement. Together, these requirements govern the breadth of claims that are granted. When the statutory patent life is the same for all patentable innovations, breadth is the main lever to differentiate rewards. Our prescription is therefore that the patent office and courts should grant generous claims (broad patents) when ideas are scarce, or more

---

[12] This result contrasts with previous treatments of secrecy in the literature. See, for example, Denicolo and Franzoni (2004) and Erkal (2005).

[13] For a sample of the many ways, other than patents, that economists have thought about incentives in R&D, see Wright (1983), chapters 2 and 8 of Scotchmer (2004) and Hopenhayn, Llobet and Mitchell (2006).

particularly, when the innovation arrives after long delay.

Patent law doctrine also has a threshold standard for granting a patent, namely, the nonobviousness requirement. Our arguments can be interpreted to mean that this threshold standard should be interpreted more leniently when ideas are scarce. In fact, patent doctrine has its own term for this circumstance, namely, "long felt need." Long-felt need is one of the secondary considerations for patentability.

7 Appendix

7.1 Proof of Lemma 1

The lemma follows from Claim 1. When the stochastic process is given by (2) as in section 3, and h is given by (7), the hypothesis of the Claim is satisfied because

$$\frac{d}{dt} h\left(\lambda | t, c\right) = F\left(c\left(t\right)\right) h\left(\lambda | t, c\right) \left[E\left(\lambda | t, c\right) - \lambda\right] \tag{22}$$

Claim 1 *Suppose there exists $\hat{\lambda}$ such that $\frac{d}{dt} h\left(\lambda | t, c\right) > 0$ for $\lambda < \hat{\lambda}$ and $\frac{d}{dt} h\left(\lambda | t, c\right) < 0$ for $\lambda > \hat{\lambda}$. Then $\frac{d}{dt} H\left(\lambda | t, c\right) > 0$ at each $\lambda \in [0, \infty)$.*

Proof: For each $\lambda \in [0, \infty)$,

$$0 = \frac{d}{dt} \int_0^\infty h\left(\tilde{\lambda} | t, c\right) d\tilde{\lambda} = \frac{d}{dt} \int_0^\lambda h\left(\tilde{\lambda} | t, c\right) d\tilde{\lambda} + \frac{d}{dt} \int_\lambda^\infty h\left(\tilde{\lambda} | t, c\right) d\tilde{\lambda}$$

For $\lambda \leq \hat{\lambda}$,

$$\frac{d}{dt} H\left(\lambda | t, c\right) = \int_0^\lambda \frac{d}{dt} h\left(\tilde{\lambda} | t, c\right) d\tilde{\lambda} > 0$$

For $\lambda > \hat{\lambda}$,

$$\begin{aligned} \frac{d}{dt} H\left(\lambda | t, c\right) &= \int_0^\lambda \frac{d}{dt} h\left(\tilde{\lambda} | t, c\right) d\tilde{\lambda} = \int_0^\infty \frac{d}{dt} h\left(\tilde{\lambda} | t, c\right) d\tilde{\lambda} - \int_\lambda^\infty \frac{d}{dt} h\left(\tilde{\lambda} | t, c\right) d\tilde{\lambda} \\ &= 0 - \int_\lambda^\infty \frac{d}{dt} h\left(\tilde{\lambda} | t, c\right) d\tilde{\lambda} > 0 \end{aligned}$$

Therefore, $H\left(\cdot | t_1, c\right)$ stochastically dominates $H\left(\cdot | t_2, c\right)$ and $(d/dt) E\left(\lambda | t, c\right) < 0$.

7.2 Proof of Proposition 1

First, the optimizing function c cannot be "U-shaped" on any domain. This implies, using (11), that $\tilde{V}$ cannot be "hill-shaped."

If the function c is "U-shaped" on some domain, there exist t_1 and t_2 such that $t_1 < t_2$, $c(t_1) = c(t_2)$, and $c'(t_1) < 0 < c'(t_2)$. However, this generates a contradiction. It holds that $\left(\frac{v}{r} - E_F(c(t_1))\right) = \left(\frac{v}{r} - E_F(c(t_2))\right)$, $\left(\frac{v}{r} - c(t_1)\right) = \left(\frac{v}{r} - c(t_2)\right)$, $F(c(t_1)) = F(c(t_2))$, and (using Lemma 1) $E(\lambda|t_1, c) > E(\lambda|t_2, c)$. Hence, using (9), this implies $\frac{d}{dt}\tilde{V}\left(t_1, c, \tilde{h}\right) < \frac{d}{dt}\tilde{V}\left(t_2, c, \tilde{h}\right)$. Together with $c'(t_1) < 0 < c'(t_2)$, this contradicts (11).

In addition, $\tilde{V}\left(t, c, \tilde{h}\right)$ is not constant on t. Using (11), that would imply that c is constant, but these two circumstances, together with the fact that $E(\lambda|t, c)$ decreases with t, contradict (9).

Thus, $\tilde{V}$ is either U-shaped or monotonic. We now show that $\tilde{V}$ is decreasing as t becomes large, hence not U-shaped, and therefore decreasing everywhere. Using (11), c is increasing.

If $\tilde{V}$ is increasing as t becomes large, the derivative $d\tilde{V}/dt$ converges to zero, since $\tilde{V}$ is bounded above by v/r. The first term on the righthand side of (10) is strictly positive as t becomes large. Thus, if the derivative converges to zero, the second term must remain strictly negative. If $\tilde{V}$ is increasing, then c is decreasing and converges to some $c^* < v/r$. If $c^* = 0$, then $F(c(t))$ converges to zero as t becomes large, which is a contradiction. If $c^* > 0$, then by the following Claim, $E(\lambda|t, c)$ converges to zero, which is also a contradiction. Hence, to complete the proof, we only need to prove the following Claim.

An intuition for the claim is that when $c(t) \to c^* > 0$, if λ is positive, then there is a positive probability of a viable idea in every time period, and an innovation will occur (the process will stop) at some finite t with probability close to one. The process can only continue to $t \to \infty$ if λ is close to 0.

Claim 2 *If $c(t) \to c^* > 0$ as $t \to \infty$, then $E(\lambda|t, c) \to 0$.*

Proof: Let $\hat{\lambda} > 0$, which also implies that $H\left(\hat{\lambda}\right) > 0$. We will show that $h\left(\hat{\lambda}|t,c\right) \to 0$.

$$
\begin{aligned}
h\left(\hat{\lambda}|\hat{t},c\right) &= \frac{\tilde{h}\left(\hat{\lambda}\right)e^{-\Lambda\left(t,c,\hat{\lambda}\right)}}{\int \tilde{h}\left(\lambda\right)e^{-\Lambda\left(t,c,\lambda\right)}d\lambda} = \frac{\tilde{h}\left(\hat{\lambda}\right)}{\int_0^\infty \tilde{h}\left(\lambda\right)e^{-\Lambda\left(t,c,\lambda\right)+\Lambda\left(t,c,\hat{\lambda}\right)}d\lambda} \\
&< \frac{\tilde{h}\left(\hat{\lambda}\right)}{\int_0^{\hat{\lambda}} \tilde{h}\left(\lambda\right)e^{-\Lambda\left(t,c,\lambda\right)+\Lambda\left(t,c,\hat{\lambda}\right)}d\lambda} = \frac{\tilde{h}\left(\hat{\lambda}\right)}{\int_0^{\hat{\lambda}} \tilde{h}\left(\lambda\right)e^{\left(\hat{\lambda}-\lambda\right)\int_0^t F\left(c\left(\hat{t}\right)\right)d\hat{t}}d\lambda} \\
&< \frac{\tilde{h}\left(\hat{\lambda}\right)}{\int_0^{\hat{\lambda}} \tilde{h}\left(\lambda\right)e^{\left(\hat{\lambda}-\lambda\right)tF\left(c^*\right)d\hat{t}}d\lambda} \to 0 \qquad \square
\end{aligned}
$$

7.3 Proof of Proposition 2

First, $\tilde{B}$ is not "hill-shaped." We show this in a slightly different way than in the proof of Proposition 1, namely, by showing that if the derivative of $\tilde{B}$ with respect to t is zero, then the second derivative of $\tilde{B}$ is positive. The first derivative is given by (15). Differentiating (15), if $d\tilde{B}\left(t,c,\tilde{h}\right)/dt = 0$, then $dc\left(t\right)/dt = d\mathcal{F}\left(t|c\right)/dt = dE_F\left(c\left(t\right)\right)/dt = 0$, and the second derivative satisfies

$$
\frac{d^2}{dt^2}\tilde{B}\left(t,c,\tilde{h}\right) = F\left(c\left(t\right)\right)\frac{d}{dt}E\left(\lambda|t,c\right)\left(\tilde{B}\left(t,c,\tilde{h}\right) - \left(\frac{v}{r} - E_F\left(c\left(t\right)\right)\right)\right) > 0
$$

The same argument shows that $\tilde{B}$ is not constant on any interval, since the second derivative of $\tilde{B}$ is not zero when the first derivative is zero.

Thus, $\tilde{B}$ is either monotonic or U-shaped. As in the proof of Proposition 1, we now show that $\tilde{B}$ cannot be increasing as t becomes large, and therefore $\tilde{B}$ is decreasing, hence the optimal threshold function c is increasing.

If $\tilde{B}$ is increasing, the derivative $d\tilde{B}/dt$ converges to zero, since $\tilde{B}$ is bounded above by v/r. The first term on the righthand side of (16) is strictly positive as t becomes large. Thus, if the derivative converges to zero, the second term must remain strictly negative. If $\tilde{B}$ is increasing, then c is decreasing and converges to some $c^* < v/r$. If $c^* = 0$, then $\mathcal{F}\left(c\left(t\right)\right)$ converges to zero as t becomes large (because $c\left(t\right) \to 0$ and $c'\left(t\right) \to 0$), which is a contradiction. If $c^* > 0$, then by Claim 2 in the proof of Proposition 1, $E\left(\lambda|t,c\right)$ converges to zero, which is also a contradiction.

27

7.4 Proof of Proposition 3

It is enough to show that $\tilde{B}\left(t,c,\tilde{h}\right) > \tilde{V}\left(t,c,\tilde{h}\right)$ for every increasing threshold function c. Define functions g, $\tilde{Q}_V$ and $\tilde{Q}_B$ by

$$g\left(\hat{t}|c\right) = e^{-r\left(\hat{t}-t\right)}\left(\frac{v}{r} - E_F\left(c\left(\hat{t}\right)\right)\right)$$

$$\tilde{Q}_V\left(\hat{t}|c\right) = \int_0^\infty \tilde{h}\left(\lambda\right) e^{-\Lambda\left(\lambda,\hat{t},c\right)}\, d\lambda$$

$$\tilde{Q}_B\left(\hat{t}|c\right) = \int_0^\infty \tilde{h}\left(\lambda\right) e^{-\lambda F\left(c\left(\hat{t}\right)\right)\hat{t}}\, d\lambda$$

$$\tilde{V}\left(t,c,\tilde{h}\right) = \int_t^\infty \int_0^\infty g\left(\hat{t}|c\right) \lambda F\left(c\left(\hat{t}\right)\right) e^{-\left[\Lambda\left(\lambda,\hat{t},c\right)-\Lambda\left(\lambda,t,c\right)\right]} h\left(\lambda|t,c\right)\, d\lambda d\hat{t}$$

$$= \frac{1}{\tilde{Q}_V\left(t|c\right)} \int_t^\infty \int_0^\infty g\left(\hat{t}|c\right) \lambda F\left(c\left(\hat{t}\right)\right) \tilde{h}\left(\lambda\right) e^{-\Lambda\left(\lambda,\hat{t},c\right)}\, d\lambda d\hat{t}$$

$$= \frac{1}{\tilde{Q}_V\left(t|c\right)} \int_t^\infty g\left(\hat{t}|c\right) F\left(c\left(\hat{t}\right)\right) \int_0^\infty \lambda\tilde{h}\left(\lambda\right) e^{-\Lambda\left(\lambda,\hat{t},c\right)}\, d\lambda d\hat{t}$$

$$= \frac{1}{\tilde{Q}_V\left(t|c\right)} \int_t^\infty g\left(\hat{t}|c\right) \frac{d}{d\hat{t}}\left[1 - \tilde{Q}_V\left(\hat{t}|c\right)\right]\, d\hat{t}$$

$$\tilde{B}\left(t,c,\tilde{h}\right) = \int_t^\infty \int_0^\infty g\left(\hat{t}|c\right) \lambda \mathcal{F}\left(\hat{t}|c\right) e^{-\lambda F\left(c\left(\hat{t}\right)\right)\hat{t}+\lambda F\left(c\left(t\right)\right)t} h\left(\lambda|t,c\right)\, d\lambda d\hat{t}$$

$$= \frac{1}{\tilde{Q}_B\left(t|c\right)} \int_t^\infty \int_0^\infty g\left(\hat{t}|c\right) \lambda \mathcal{F}\left(\hat{t}|c\right) \tilde{h}\left(\lambda\right) e^{-\lambda F\left(c\left(\hat{t}\right)\right)\hat{t}}\, d\lambda d\hat{t}$$

$$= \frac{1}{\tilde{Q}_B\left(t|c\right)} \int_t^\infty g\left(\hat{t}|c\right) \mathcal{F}\left(\hat{t}|c\right) \int_0^\infty \lambda\tilde{h}\left(\lambda\right) e^{-\lambda F\left(c\left(\hat{t}\right)\right)\hat{t}}\, d\lambda d\hat{t}$$

$$= \frac{1}{\tilde{Q}_B\left(t|c\right)} \int_t^\infty g\left(\hat{t}|c\right) \frac{d}{d\hat{t}}\left[1 - \tilde{Q}_B\left(\hat{t}|c\right)\right]\, d\hat{t}$$

At each $\hat{t} > 0$, $e^{-\Lambda\left(\lambda,\hat{t},c\right)} > e^{-\lambda F\left(c\left(\hat{t}\right)\right)\hat{t}}$. Therefore,

$$\tilde{Q}_V\left(\hat{t}|c\right) = \int_0^\infty \lambda\tilde{h}\left(\lambda\right) e^{-\Lambda\left(\lambda,\hat{t},c\right)} d\lambda > \int_0^\infty \lambda\tilde{h}\left(\lambda\right) e^{-\lambda F\left(c\left(\hat{t}\right)\right)\hat{t}} d\lambda = \tilde{Q}_B\left(\hat{t}|c\right)$$

The functions $1 - \tilde{Q}_V\left(\cdot|c\right)$ and $1 - \tilde{Q}_B\left(\cdot|c\right)$ are distribution functions, and $1 - \tilde{Q}_V\left(\cdot|c\right)$ stochastically dominates $1-\tilde{Q}_B\left(\cdot|c\right)$. Since $g\left(\cdot\right)$ is decreasing when c is increasing, it follows that $\tilde{Q}_B\left(t|c\right)\tilde{B}\left(t,c,\tilde{h}\right) > \tilde{Q}_V\left(t|c\right)\tilde{V}\left(t,c,\tilde{h}\right)$.

But since $\tilde{Q}_B\left(t|c\right)/\tilde{Q}_V\left(t|c\right) < 1$, it follows that $\tilde{B}\left(t,c,\tilde{h}\right) > \tilde{V}\left(t,c,\tilde{h}\right)$. $\square$

7.5 Proof of Remark 2

(a) Let n be the number of potential recipients. Suppose that each recipient believes that each other recipient follows an investment strategy described by a threshold function c. At date t, some recipients have received ideas. If a recipient received a single idea at, say $\tilde{t} \leq t$, the recipient takes this into account in forming his belief on λ. In a population of size n, with individual arrival rate λ/n, the agent's belief on λ is given by the following posterior density:

$$g_n\left(\hat{\lambda}|\tilde{t}, t, c\right) = \frac{\frac{\hat{\lambda}}{n} e^{-\frac{\hat{\lambda}}{n}\tilde{t}} e^{-\left(\frac{n-1}{n}\right)\hat{\lambda}F(c(t))t}\tilde{h}\left(\hat{\lambda}\right)}{\int \frac{\lambda}{n} e^{-\frac{\lambda}{n}\tilde{t}} e^{-\left(\frac{n-1}{n}\right)\lambda F(c(t))t}\tilde{h}\left(\lambda\right)d\lambda} = \frac{\hat{\lambda}e^{-\left(\frac{n-1}{n}\right)\hat{\lambda}F(c(t))t}\tilde{h}\left(\hat{\lambda}\right)}{\int \lambda e^{-\frac{(\lambda-\hat{\lambda})}{n}\tilde{t}} e^{-\left(\frac{n-1}{n}\right)\lambda F(c(t))t}\tilde{h}\left(\lambda\right)d\lambda}$$

Let $\hat{h}$ be the limit density function as $n \to \infty$:

$$\hat{h}\left(\hat{\lambda}|t, c\right) = \frac{\hat{\lambda}e^{-\hat{\lambda}F(c(t))t}\tilde{h}\left(\hat{\lambda}\right)}{\int \lambda e^{-\lambda F(c(t))t}\tilde{h}\left(\lambda\right)d\lambda}$$

The limit distribution does not depend on $\tilde{t}$, as asserted in part (a).

Further, the probability of receiving more than one idea has a second-order effect, and we therefore ignore it. The numerator in the following expression is the probability of receiving two or more ideas by time t, and the denominator is the probability of receiving one or more ideas by time t. Using L'Hopital's rule, the ratio converges to zero.

$$\lim_{n\to\infty} \frac{\left[1 - e^{-\frac{\lambda}{n}t} - \left(\frac{\lambda}{n}t\right)e^{-\frac{\lambda}{n}t}\right]}{\left[1 - e^{-\frac{\lambda}{n}t}\right]} \to 0$$

This concludes part (a).

(b) Nevertheless, the recipients are more optimistic about λ than the planner. When the planner believes that c describes the recipients' investment behavior, the density of the planner's posterior is

$$h\left(\hat{\lambda}|t, c\right) = \frac{e^{-\hat{\lambda}F(c(t))t}\tilde{h}\left(\hat{\lambda}\right)}{\int e^{-\lambda F(c(t))t}\tilde{h}\left(\lambda\right)d\lambda}$$

The ratio of the densities, $\hat{h}\left(\hat{\lambda}|t, c\right)/h\left(\hat{\lambda}|t, c\right)$, is proportional to $\hat{\lambda}$. This implies that the recipients, as opposed to the planner, place higher weight on higher λ, that $\hat{h}$ stochastically

dominates h, and that $\hat{h}$ has a higher expected value. It is instructive to show the latter directly.

Divide both numerator and denominator of $\hat{h}$ by $\int e^{-\lambda F(c(t))t}\tilde{h}(\lambda)\,d\lambda$:

$$\hat{h}\left(\hat{\lambda}|t,c\right)=\frac{\left(\frac{\hat{\lambda}e^{-\hat{\lambda}F(c(t))t}\tilde{h}(\hat{\lambda})}{\int e^{-\lambda F(c(t))t}\tilde{h}(\lambda)d\lambda}\right)}{\left(\frac{\int \lambda e^{-\lambda F(c(t))t}\tilde{h}(\lambda)d\lambda}{\int e^{-\lambda F(c(t))t}\tilde{h}(\lambda)d\lambda}\right)}=\frac{\hat{\lambda}e^{-\hat{\lambda}F(c(t))t}\tilde{h}\left(\hat{\lambda}\right)}{E\left(\lambda|t,c\right)\int e^{-\lambda F(c(t))t}\tilde{h}\left(\lambda\right)d\lambda}\tag{23}$$

Let $E\left(\lambda|t,c\right)$ and $E\left(\lambda^2|t,c\right)$ be the expected values with respect to the planner's posterior belief, and let $Var\left(\lambda\right)$ be the variance of the planner's belief. Now consider the expected value of λ with respect to the recipients' belief instead of the planner's belief. Using (23),

$$
\begin{aligned}
\hat{E}\left(\lambda|t,c\right) &= \int \hat{\lambda}\hat{h}\left(\hat{\lambda}|t,c\right)d\hat{\lambda}=\frac{\int \hat{\lambda}^2\left(\frac{e^{-\hat{\lambda}F(c(t))t}\tilde{h}(\hat{\lambda})}{\int e^{-\lambda F(c(t))t}\tilde{h}(\lambda)d\lambda}\right)d\hat{\lambda}}{E\left(\lambda|t,c\right)}=\frac{E\left(\lambda^2|t,c\right)}{E\left(\lambda|t,c\right)}\\
&= \left(\frac{1}{E\left(\lambda|t,c\right)}\right)\left(Var\left(\lambda\right)+E\left(\lambda|t,c\right)^2\right)=\frac{Var\left(\lambda\right)}{E\left(\lambda|t,c\right)}+E\left(\lambda|t,c\right)
\end{aligned}
$$

Thus, $\hat{E}\left(\lambda|t,c\right)>E\left(\lambda|t,c\right)$. $\square$

References

[1] Aghion, P. and P. Howitt. 1992. "A Model of Growth Through Creative Destruction," *Econometrica*, 60(2), 323-351.

[2] Berliant, M. and M. Fujita. 2008. "Knowledge Creation as a Square Dance on the Hilbert Cube," *International Economic Review,* 49, 1251-1295.

[3] Choi, J. P. 1991. "Dynamic R&D Competition under 'Hazard Rate' Uncertainty," *RAND Journal of Economics,* 22, 596-610

[4] Denicolo, V. and L. Franzoni. 2004. "Patents, Secrets, and the First Inventor Defense," *Journal of Economics and Management Strategy*, 13, 517-538.

[5] Dixit, A. and R. Pindyck. 1994. *Investment under Uncertainty.* Princeton, NJ: Princeton University Press.

[6] Erkal, N. 2005. "The Decision to Patent, Cumulative Innovation, and Optimal Policy," *International Journal of Industrial Organization*, 23(7-8), 535-562.

[7] Erkal, N. and S. Scotchmer. 2007. "Scarcity of Ideas and Options to Invest in R&D," University of California, Berkeley, Department of Economics, Working Paper 07-348.

[8] Erkal, N. and S. Scotchmer. 2009. "Scarcity of Ideas and R&D Options: Use it, Lose it, or Bank it," NBER Working Paper No. 14940.

[9] Hopenhayn, H., G. Llobet, and M. Mitchell. 2006. "Rewarding Sequential Innovators: Patents, Prizes, and Buyouts." *Journal of Political Economy* 114:1041-1068.

[10] McCall, B. P. and J. J. McCall. 2008. *Economics of Search.* New York: Routledge Press.

[11] MacDonald, R. and D. Siegel. 1986. "The Value of Waiting to Invest," *Quarterly Journal of Economics*, 101, 707-728.

[12] Malueg, D. A. and S. O. Tsutsui. 1997. "Dynamic R&D Competition with Learning," *RAND Journal of Economics* 28, 751-772.

[13] O'Donoghue, T., Scotchmer, S. and Thisse, J.-F. 1998. "Patent Breadth, Patent Life and the Pace of Technological Progress," *Journal of Economics and Management Strategy*, 7, 1-32.

[14] Reinganum, J. 1989. "The Timing of Innovation: Research, Development and Diffusion," in R. Schmalensee and R. D. Willig, eds., *Handbook of Industrial Organization*, Amsterdam: Elsevier.

[15] Rosenfield, D. B. and R. D. Shapiro. 1981. "Optimal Adaptive Price Search," *Journal of Economic Theory*, 25, 1-20.

[16] Rothschild, M. 1974. "Searching for the Lowest Price when the Distribution of Prices is Unknown," *Journal of Political Economy*, 82, 689-711.

[17] Scotchmer, S. 1999. "On the Optimality of the Patent Renewal System," *RAND Journal of Economics,* 30, 131-196.

[18] Scotchmer, S. 2004. *Innovation and Incentives.* Cambridge, MA: MIT Press.

[19] Snyder, D. L. and M. I. Miller. 1991. *Random Point Processes in Time and Space.* New York: Springer-Verlag.

[20] Weeds, H. 2002. "Strategic Delay in a Real Options Model of R&D Competition," *Review of Economic Studies*, 69, 729-747.

[21] Wright, B. D. 1983. "The Economics of Invention Incentives: Patents, Prizes and Research Contracts." *American Economic Review* 73:691-707.