

Garfagnini, Umberto; Strulovici, Bruno

Working Paper

Social learning and innovation cycles

CSIO Working Paper, No. 0114

Provided in Cooperation with:

Department of Economics - Center for the Study of Industrial Organization (CSIO), Northwestern University

Suggested Citation: Garfagnini, Umberto; Strulovici, Bruno (2012) : Social learning and innovation cycles, CSIO Working Paper, No. 0114, Northwestern University, Center for the Study of Industrial Organization (CSIO), Evanston, IL

This Version is available at:

<https://hdl.handle.net/10419/92479>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

THE CENTER FOR THE STUDY
OF INDUSTRIAL ORGANIZATION
AT NORTHWESTERN UNIVERSITY

Working Paper #0114

Social Learning and Innovation Cycles

By

Umberto Garfagnini
ITAM School of Business

and

Bruno Strulovici
Northwestern University

February 6, 2012

Social Learning and Innovation Cycles

Umberto Garfagnini Bruno Strulovici
ITAM School of Business Northwestern University

February 6, 2012

Abstract

We study social learning and innovation in an overlapping generations model, emphasizing the trade-off between marginal innovation (combining existing technologies) and radical innovation (breaking new ground). We characterize both short-term and long-term dynamics of innovation, and the intergenerational accumulation of knowledge. Innovation cycles emerge endogenously, but the number of cycles is finite almost surely, and radical innovation terminates in finite time. We identify a negative relationship between past successes and the magnitude of radical innovation, combining insights from the multi-armed bandit literature with a spatial representation of innovation. Past successes reduce the incremental value of experimentation, and result in less ambitious innovation. In our framework, patents promote radical innovation through two channels: by increasing the expected benefit of radical innovation and by increasing the cost of marginal innovation. Our analysis suggests that sustaining radical innovation in the long-run requires external intervention.

This paper was previously circulated under the title “The Dynamics of Innovation”. We are grateful to Alessandro Bonatti, Steve Callander, Xavier Durand, Michael Mandler, Kiminori Matsuyama, Joel Mokyr, Francisco Ruiz-Aliseda, Simon Wilkie, Huanxing Yang, and seminar participants at the Spring 2011 Midwest Theory Conference, Banco de México, and the 2012 North American Winter Meeting of the Econometric Society. Email addresses: umberto.garfagnini@itam.mx and b-strulovici@northwestern.edu.

1 Introduction

The invention of the transistor in the 1950s, the first large-scale cellular-phone system in the 1970s, and the internet in the 1990s are all radical historical breakthroughs that boosted the development of new industries (e.g., semiconductors, computers, e-commerce, etc.), which shaped the world as we know it. The importance of radical innovation from a social perspective has ample empirical evidence (Griliches (1992), Hall (1996)), but is such innovation sustainable in the long-run?

In the last few decades, many research institutions, such as Bell Labs and the Defense Advanced Research Projects Agency, have shifted funding from basic to more applied research. Those institutions contributed to the economic boom of the US economy with their innovative research in the second half of the twentieth century.¹ In a recent paper, Jones and Williams (1998) conclude that “the optimal share of resources to invest in research is conservatively estimated to be two to four times larger than the actual amount invested by the U.S. economy. The extent of underinvestment is substantial, and could well be much larger.” While this behavior may be optimal in the short-term because it leads to higher profitability, it undermines long-run growth.

This paper studies the short-run and long-run dynamics of innovation in a model of social learning where technologies have correlated payoffs. Adopting an “experimentation” approach, we focus on the dynamics of innovation and knowledge, rather than on their implications for capital accumulation, which distinguishes our work from models in which innovation takes a more reduced form.² At each period, a “young” agent (or generation) can experiment with a technology chosen on the positive real half-line, whose payoff is initially unknown and drawn from the path of an arithmetic Brownian motion. Brownian motion is the only continuous process (up to scaling and drift parameters) that satisfies some stationary, independent increments condition,³ and is thus a natural way of modeling unknown but

¹We refer to the article “*How Science Can Create Millions of New Jobs: Reigniting basic research can repair the broken U.S. business model and put Americans back to work*”, published in BusinessWeek on August 27th, 2009, for a discussion of this issue.

²In our model, aggregate shocks arise at each period, and their distribution depends on the type of innovation that is chosen at that period, as well as on past innovation. The relation to growth is discussed in Section 8.

³This result directly follows from Lévy’s characterization theorem for the representation of continuous

correlated technologies.⁴ Each agent lives for two periods and, therefore, has an incentive to experiment in its first period. We distinguish two types of innovation. *Marginal innovation* consists in choosing a technology in the *explored hull*, i.e., a technology that is a convex combination of technologies that have been explored earlier. Marginal innovation bears a cost that is normalized to zero.⁵ *Radical innovation* consists in experimenting beyond the frontier of that interval, and entails a cost that is increasing and convex in the distance from the frontier, and may decrease with the quality of the best explored technology.

The paper’s first contribution is a representation of the *value of marginal innovation* and the *value of radical innovation*, defined, at each period, as the expected benefit of each type of innovation in excess of the payoff of the best technology explored so far. Using those values of innovation, we are able to precisely describe the dynamics of innovation.

We identify a negative relationship between the magnitude of radical innovation (i.e., how far it is from known technologies) and past successes. This effect stems from the relation between the “value of experimentation,” a well-known concept in the experimentation literature, with our spatial representation of technologies. More precisely, past successes reduce not only the value of experimentation, but also its *derivative* with respect to the size of innovation: it reduces the *marginal* value of experimentation where the term “marginal” refers to incremental departures from known technologies (the spatial dimension of the model). Because the marginal cost of radical innovation is nondecreasing in the distance from known technologies, the optimal size of radical innovation must therefore decrease with past successes. By a similar argument, more patient agents undertake more ambitious radical innovation, because the marginal value of experimentation lies in the future exploitation of information acquired today, whereas the cost of innovation is incurred in the present period.

We also characterize *innovation cycles*: successive agents alternate between radical and marginal innovation, in a pattern reminiscent of Schumpeterian cycles. If radical innovation at some period yields a high enough payoff, it is followed by further radical innovation.

martingales (see, e.g., Karatzas and Shreve 1991).

⁴A consequence of this feature is that the payoff space is unbounded, a characteristic that requires a significantly more sophisticated analysis for the long-run behavior of innovation, but which gives a strong meaning to our result on stagnation (Theorem 2).

⁵Our results are qualitatively unchanged if the cost of marginal innovation is low, compared to the cost radical innovation.

Otherwise, the refinement of known technologies takes place through marginal innovation. However, marginal innovation reduces the value of experimentation inside the explored hull, other things equal. Thus, after a sequence of marginal innovations, the value of marginal innovation can drop below the value of radical innovation, and radical innovation becomes again optimal, triggering a new innovation cycle.

Again, this phenomenon stems from the interplay between a well-known mechanism in the experimentation literature and the spatial nature of our model. In a multi-armed bandit problem, the more an “arm” is played, and the lower the “pure” value of experimentation for that arm: by playing this arm, one reduces the uncertainty surrounding the payoff distribution of that arm, and this reduces the learning value of playing that arm. Marginal innovation has a similar effect: by trying repeatedly convex combinations of known technologies, one reduces uncertainty about *all* technologies in the explored hull, due to the correlation between these technologies. By contrast, marginal innovation does not affect beliefs about radical innovation, and thus does not reduce the experimentation value of radical innovation.

Formalizing this intuition requires a new relation between “arms” and “technologies,” which is another contribution of our paper. Because Brownian motion is a Markov process, the payoff distribution of all technologies within any interval spanned by two technologies with known payoffs is *conditionally independent* of all technologies located outside of that interval. Therefore, the standard intuition for the value of experimentation can be applied at the level of these intervals, or “units,” as we call them, rather than at the lower level of individual technologies. The more technologies are tried within a given unit, and the lower the experimentation value of that unit.

We exploit this representation of the technological space to derive an “index” policy for experimentation. Each unit is assigned an index which has an exact formula up to a function that depends only on the ratio between the width of the unit and the difference of payoffs at its endpoints. The wider the unit, and the higher its value of experimentation. Because an “old” agent always picks the best technology explored so far, the index of a unit also depends on the payoff of that technology:⁶ a better known technology lowers, all else equal, the prob-

⁶This distinguishes this index from the well-known Gittins index, whose value for an arm is completely independent of other arms.

ability of surpassing it through experimentation, and lowers the index of all units. However, conditional independence implies that, other than through the best available technology, the index of a unit is completely independent of the shape and level of other units.

We show that radical innovation stops in finite time and, therefore, that the number of innovation cycles is (almost surely) finite. Surprisingly, this stagnation arises even when the underlying Brownian motion has any positive drift, as long as the decision problem of each generation is well defined (see Section 7), and may be explained as follows. Each radical innovation has a positive probability of being a large failure, independently of the particular experimentation path. After such failure, the next generation finds it optimal to perform marginal innovation, rather than undertake a radical innovation building on such a failure. Even with a positive drift, the large leap one would have to make to offset this failure entails such a high innovation cost that radical innovation is suboptimal. This leads to a vicious cycle: because the first generation following the failure does not try radical innovation, the future generations' perspective of radical innovation does not improve, and because it does not improve, those later generations do not try radical innovation either. If some generation internalized the benefit of radical innovation for future generations, it could break this cycle. Short of this long-sighted perspective, radical innovation stops inefficiently early.

Once the “stagnation” stage has been reached, new generations fine-tune their search to make marginal improvements over available technologies. Since opportunities for valuable innovation become increasingly rare in a stagnating economy, the “value” of pursuing marginal innovation converges to zero in the limit. Innovations then converge to a single technology, which can be interpreted as the emergence of a technological standard. Another interpretation of this result is that the economy converges to a balanced-growth path, when we reinterpret the outcome of each technology as a productivity rate. In our model, therefore, social learning becomes arbitrarily small. In contrast to the usual mechanism of informational cascades of Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992), the reason here is that information acquisition is endogenous, and becomes asymptotically negligible. The mechanism through which radical innovation stops shares some features with information cascades however: because past generations do not learn anything, after some time, about radical innovation, the value of radical innovation stalls, and future generations have no incentive to experiment with radical innovation, creating a vicious cycle.

Characterizing the limit of beliefs, and innovation in our setting presents a technical challenge. For example, could agents stop experimentation altogether, and start exploiting the best known technology after some time? Or could they switch infinitely often between different “technological areas”? Limiting beliefs result from the interaction between knowledge and decisions at each time. The challenge, in our setting, stems from the combination of three elements: i) a continuum of actions, ii) correlation across the payoffs of all actions, and iii) non compactness of the action space. The non compactness is an integral part of our stagnation result, and gives it its strength. However, it requires specific arguments to adapt the standard measure theoretical tools and results developed for compact spaces. Specifically, one contribution of this paper is to establish the almost sure convergence of these beliefs (and, therefore, of social learning) to a well-defined distribution over the set of technological payoffs, in a setting with Brownian uncertainty.

Our model gives rise to multiple types of intergenerational linkages: A direct *learning effect*, which arises from the observability of the outcomes of innovation by previous generations. Innovation increases the stock of knowledge of society through more refined beliefs about the underlying outcome process. A *feasibility effect*, which arises when radical innovation is undertaken, because it endows society with new opportunities for costless marginal innovation. A third intergenerational linkage is the *cost effect*, which arises when the cost of radical innovation is reduced following new discoveries, and is explored in Section 7. Feasibility and cost effects strengthen the incentives for radical innovation. While the economy still experiences stagnation in the long-run, the short-run dynamics may be significantly affected by such effects.

Our paper is closely related to Jovanovic and Rob (1990), who study a search-theoretic model of growth through technological innovation. One major departure from that paper is the endogenous relation between incentives for future innovation and current technologies. In our model, the incentives to perform radical innovation are highly sensitive to the history of previous discoveries in two ways. First, there is a payoff externality that derives from the possibility of exploitation of previous technologies, as in Jovanovic and Rob. Second, the history of technologies directly affects the beliefs that each agent holds about the outcome of radical innovation, which is absent in their paper. Our model generates innovation cycles, as in Jovanovic and Rob (1990). In our model, the size of radical innovation is determined

endogenously, affecting the length of each cycle.⁷

The paper is also related to the model introduced by Callander (2011), who analyzes the experimentation path produced by a sequence of myopic agents trying to find the “right” policy. Because, in our model, agents live for two periods, learning is *active* rather than *passive*. Active learning drives incentives to innovate with new technologies through the option value of experimentation. Our model also distinguishes between radical and marginal innovation, giving rise to innovation cycles based on the fluctuating option value for innovation. Our stagnation and convergence results are also of a very different nature from the long-run dynamics arising in Callander’s model.⁸

We also contribute to the literature on optimal experimentation following Rothschild (1974), McLennan (1984), Easley and Kiefer (1988), Aghion, Bolton, Harris, and Jullien (1991), and Bala and Goyal (1998), among others. The literature has focused on the possibility of incomplete learning in the long-run and has tried to identify the conditions under which learning can be guaranteed to be complete. In contrast, learning can hardly be expected to be complete in our set-up, as the underlying parameter is the realized path of a Brownian motion. Some of our arguments build on Easley and Kiefer (1988), who derive the properties of the experimentation policy in the long-run. In our setting, not only does there exist a well-defined long-run belief about the value of each policy: we show that experimentation converges to a *single* policy. Exploiting the properties of Brownian motion, we prove our results without assuming compactness of the parameter space, in contrast to Easley and

⁷Our model is also related to the literature on organizational behavior, which stresses the tendency of firms to perform local searches around the technologies currently in use. This is a natural response to uncertainty about the outcomes of innovation. For example, Kauffman, Lobo, and Macready (2000) look at environments in which a firm’s current location in the space of technologies influences the incentives for innovation. In a different framework, Matsuyama (1999) investigates the connection between neoclassical and neo-Schumpeterian growth models. Acemoglu, Gancia, and Zilibotti (2010) study the interplay between innovation and standardization in a dynamic general equilibrium model. In contrast, we analyze innovation and standardization (i.e., radical versus marginal innovation) in a framework affected by uncertainty about the payoff associated with a new technology.

⁸A parallel line of research has constructed models of scientific research. Bramoullé and Saint-Paul (2010) studies the evolution of fundamental research when scientists are motivated by the reputation they obtain from the development of new fields. Reputational concerns increase incentives to pursue more innovative research. Using a different model, Mandler (2011) shows that reduced communication among scientists may be beneficial for society by inducing different scientists to pursue different lines of research.

Kiefer. Also in contrast to some of this literature, our paper characterizes both the short-term patterns of innovation, as well as its behavior in the long run.

Finally, the paper can also be interpreted as a theory of economic stagnation. Jovanovic and Nyarko (1996) have shown that the development of human capital specific to a particular technology may induce a myopic agent to stick to that technology in the long run, despite the availability of (possibly) better technologies. Garicano and Rossi-Hansberg (2009) provides (among other things) an alternative theory of stagnation, which is based on the necessity to develop organizations to exploit newly discovered technologies. In this paper, stagnation follows endogenously from the outcomes of the innovation process.

2 The Model

We consider an overlapping generations model with the following characteristics. An agent is born at each period t , who lives for two periods, “young” and “old.” The agent is risk neutral and chooses at each period a technology $x \in X = [0, \infty)$. Each technology x has a payoff $f(x) \in \mathbb{R}$ that is initially unknown except at the origin: $f(0) = 0$.⁹ The agent must incur the payoff of any technology he has chosen.¹⁰

A young generation inherits from the contemporary old generation the knowledge of all technologies and payoffs that have previously been tried.¹¹ A time- t history h_t is a list of technology-payoff pairs that have been experienced in the past. Letting X_t denote the set of previously tried technologies, $h_t = \{(0, 0), (x, f(x)) : x \in X_t\}$. The following variables will play an important role in the analysis. Given any history h_t , we let

- $z_t = \max\{0, f(x) : x \in X_t\}$ denote the highest payoff among explored technologies;
- $\bar{x}_t = \max\{x : x \in X_t\}$ denote the (right) *frontier* of the set of explored technologies;

⁹Instead of a payoff, $f(x)$ could represent a utility level, without affecting the analysis as long as agents have homogeneous utility and are expected-utility maximizers. One could also consider the larger domain $X = \mathbb{R}$. Radical innovation would then have “left” and “right” components.

¹⁰An alternative modeling choice would be to follow Jovanovic and Rob (1990), in which an agent can learn at some cost the payoff of a new technology, and then decide whether or not to use it.

¹¹Old and young generations coexisting in a period interact only through the transmission of knowledge of previously tried technologies. This transmission is assumed to be perfect. The two generations choose their technology simultaneously and thus cannot observe the outcome of each other’s innovation.

- $\rho_t = z_t - f(\bar{x}_t)$ denote the *gap* in values between the best explored technology and the frontier technology.

Because agents perfectly know past innovations and their outcomes, z_t is nondecreasing over time, and is a lower bound on the payoff that agents can guarantee themselves, at any period. The set $[0, \bar{x}_t]$ of technologies that are convex combinations of previously explored ones is called the *explored hull*. The gap measures the profitability differential between the best technology explored so far and the current frontier.

Each generation shares the common belief that the underlying outcome function is the realized path of a Brownian motion with drift κ and volatility parameter $\sigma > 0$. More precisely, we start with a filtered probability space $(\Omega, \mathcal{F}, \mathbb{P})$ satisfying the usual regularity conditions,¹² and whose outcomes are identified with the paths of some Brownian motion.¹³ We first focus on the case in which $\kappa = 0$ and differ to Section 7.2 the extension of our result to the case of a nonzero drift. In particular, the payoff of a technology $x > \bar{x}_t$ has a normal distribution, with the following parameters:

$$f(x) \sim \mathcal{N}(f(\bar{x}_t), \sigma^2(x - \bar{x}_t)). \quad (1)$$

Technologies beyond \bar{x}_t thus have the same expected payoff, and a variance that increases with their distance to \bar{x}_t .

To exploit the Markov property of Brownian motion, we represent past innovations in terms of technological “units,” defined as follows.

DEFINITION 1 *A bounded unit $u = [I, m, M]$ consists of a finite interval $I = [x_l, x_r]$ with ordered endpoint values $m = \min\{f(x_l), f(x_r)\}$ and $M = \max\{f(x_l), f(x_r)\}$. An unbounded unit $u = [I, m]$ consists of an interval $I = [x_l, +\infty)$ with initial value $m = f(x_l)$.*

For any bounded unit u , we let $l = x_r - x_l$ and $d = M - m$ denote the width and height of the unit. After any history, there is a single unbounded unit, which we simply denote u_∞ . We denote by $\mathcal{P}(h_t)$ the collection of *bounded* units induced by an arbitrary history h_t .

¹²See Karatzas and Shreve (1991).

¹³In addition to Lévy’s characterization theorem, mentioned in Footnote 3, Jovanovic and Rob (1990) provide an axiomatic foundation for the Brownian structure of innovations.

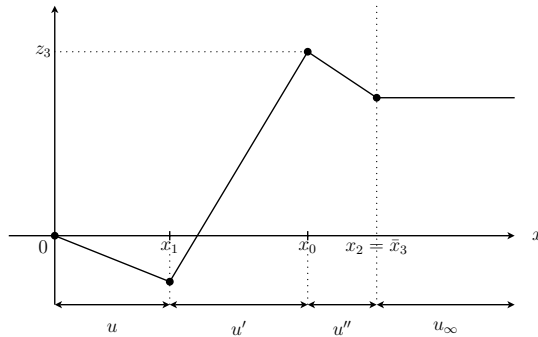


Figure 1: Example of collection of units after three periods.

A technology x in a bounded unit u has a normally distributed payoff with mean

$$f(x_l) + \frac{f(x_r) - f(x_l)}{x_r - x_l}(x - x_l) \quad (2)$$

and variance

$$\frac{(x - x_l)(x_r - x)}{x_r - x_l} \sigma^2. \quad (3)$$

The stochastic payoff function of technologies within a given unit is a Brownian bridge.¹⁴ The expected outcome of technologies within a unit increases linearly from the endpoint technology with the worst outcome to the one with the highest outcome. The variance, instead, increases as we move away from either endpoint and it is maximized at the midpoint technology. Observing the payoff of a technology in a given unit affects only the distribution of technologies lying in that unit: payoffs of technologies outside of that unit are conditionally independent from the payoffs in the unit, given the payoffs at the endpoints of the unit. Figure 1 provides an example of a collection of units where each technology is represented with the associated expected outcome.

We consider two types of innovation. A *marginal innovation* is a technological choice $x \in [0, \bar{x}_t]$, and is assumed to be costless (leaving aside the opportunity cost of forgoing the highest known payoff z_t). Thus, whenever the frontier expands, all technologies between the old and new frontiers become available at no cost.¹⁵ *Radical innovation*, instead, incurs a

¹⁴We refer the reader to Billingsley (1968) for an introduction to Brownian bridges.

¹⁵The qualitative results of the paper are robust to the introduction of a positive cost, provided that either i) exploitation incurs the same cost, or ii) the cost of marginal innovation vanishes as the innovation becomes arbitrarily close to known technologies.

cost that depends on how far an agent pushes innovation away from the current frontier. We motivate this assumption by the fact that large initial investments are arguably one of the main features of fundamental research, together with its high uncertainty. Specifically, we assume for now that the cost for innovation $x > \bar{x}$ is $c(x - \bar{x})$ for some function c .

ASSUMPTION 1 $c(\cdot)$ is twice continuously differentiable, strictly increasing, weakly convex, and such that *i*) $c(0) = 0$, and *ii*) either $c'(0) > 0$ or $c''(0) > 0$.

Section 7.1 analyzes environments in which the cost of radical innovation is also a (decreasing) function of the quality of the best explored technology.

Each agent maximizes his total expected payoff, discounting his second period payoff by a factor $\delta \in (0, 1)$. An old agent has no value for information and thus always chooses the best known technologies. A young agent solves the optimization problem

$$U(h_t) = \sup_{x \in X} E_{h_t} [f(x) - c(x - \bar{x}_t) + \delta \max\{f(x), z_t\}] \quad (4)$$

with the convention that $c(y) = 0$ whenever $y \leq 0$.

3 How to Choose Innovation

This section characterizes the equilibrium path of innovation, resulting from the optimal experimentation of each generation. The problem faced by a young agent resembles that of a multi-armed bandit problem, with important differences. Innovating with a new technology creates an informational spillover due to the correlation across technologies. A more structured way to think about the decision problem faced by the agent is to bundle technologies together according to the units they belong to. Since the experimentation of a technology only changes the payoff distribution of technologies that reside in the same unit, the agent's problem may be decomposed as, first, choosing one of finitely many units and, second, which technology to pick within that unit. For the unit-level choice, Theorem 1 shows that we can characterize each unit based on a single index. This index determines the *value of innovation* of the unit, defined as the largest (normalized) discounted expected payoff attainable by the current generation, net of the highest payoff among explored technologies, z , when the innovations are restricted to that unit in the first period, and given the possibility of exploitation with any known technology in the second period.

To state Theorem 1, we fix a history h and let z denote the highest outcome among explored technologies, $\rho = z - f(\bar{x})$ denote the gap, and \bar{x} denote the frontier. For any unit u , we recall that l denotes the width of the unit, $m \leq M$ denote the payoffs at its endpoints, and $d = M - m$.

THEOREM 1 *To each unit u corresponds an index, $\gamma(u, z)$, with the following properties:*

i) It is optimal for the agent to pick the unit with the highest index γ .

ii) The index is characterized by functions $\eta : \mathbb{R}_+^2 \rightarrow \mathbb{R}$, and $\psi : \mathbb{R}_+ \rightarrow \mathbb{R}$ such that

$$\gamma(u, z) = \begin{cases} m + d \eta\left(\frac{\sqrt{l}}{d}, \frac{z-m}{d}\right) & \text{if } u \neq u_\infty \\ f(\bar{x}) + \psi(\rho) & \text{if } u = u_\infty \end{cases} \quad (5)$$

iii) For a bounded unit, the index is increasing in l and d .

*Moreover, it is strictly optimal, for a bounded unit, to choose a technology closer to the endpoint with the higher payoff.*¹⁶

The highest known payoff z enters the index of a unit because it affects the agent's continuation payoff. Indeed, the only reason for the agent to explore a new technology in his first period is the positive probability that the payoff of this technology will surpass z .

Given a history h_t and a unit u implied by that history, let $U(u; h_t)$ denote the highest lifetime expected payoff that the young generation can get, given history h_t , when his choice in the first period is restricted to a technology within the unit u (only). We define the *value of innovation of u* as

$$V(u, z) = \frac{U(u; h_t)}{1 + \delta} - z. \quad (6)$$

The value of innovation of a unit is the per-period expected payoff that an agent get from innovating in that unit only (when young) and then exploit the best explored technology (when old), above and beyond the payoff of the best explored technology. By definition, it is suboptimal to choose any technology in a unit with a negative value of innovation, while

¹⁶As a consequence, it is optimal to choose the midpoint of a unit whose endpoints have the same payoff.

a unit with a strictly positive value of innovation dominates exploitation of the best (and, therefore, any) explored technology.

The value of innovation of a unit is closely related to the index introduced by Theorem 1: Equation (12) in the Appendix implies that

$$V(u, z) = \gamma(u, z) - z.$$

PROPOSITION 1 (VALUE MONOTONICITY) *For any bounded unit u , $V(u, z)$ is strictly decreasing in z . If, at any time, $V(u, z_t) < 0$, no technology in u is ever chosen after time t .*

To understand Proposition 1, consider, first, the simpler case of a unit u that does not contain the best explored technology, and suppose that the payoff of that technology is increased from z to $z' > z$, keeping all else constant. That increase has no effect on the payoff distribution of technologies inside u , and reduces the probability that any technology in u beats the best known technology. Thus, $V(u, z') < V(u, z)$. Now, consider a unit u whose endpoints include the best explored technology. For such unit, a higher value of z increases, linearly, the expected payoff of all technologies inside the unit. Since the variance over that unit is unaffected, however, exploitation with the best explored technology is relatively more appealing than before the increase.¹⁷

Following the experimentation literature, we call *exploitation* the fact of choosing an explored technology (and whose payoff is, therefore, perfectly known). In contrast to standard bandit models with a “safe” arm, exploitation here is strictly suboptimal at all times.

PROPOSITION 2 (EXPLOITATION) *For any unit u containing the best explored technology, $V(u, z) > 0$. In particular, exploitation is strictly suboptimal after all histories.*

The first part of the proposition is proved in the appendix. Intuitively, slightly departing from the best explored technology slightly reduces the expected payoff of the agent, but also

¹⁷A starker intuition for this result can be obtained by appealing to the theory of large deviations (see Dembo and Zeitouni (1998)): as z gets arbitrarily large, the payoff distribution inside the unit u looks closer to a straight line, joining the low-payoff extremity x_l to the best technology x_h with payoff z : $(f(x) - f(x_l))/(z - f(x_l)) \rightarrow_{z \rightarrow \infty} m + (z - m)(x - x_l)/(x_h - x_l)$ a.s., where m is the payoff at x_l . As z gets arbitrarily large, therefore, the probability that any given technology x in u surpasses z , converges to zero. Exploitation of x_h remains suboptimal for all values of z , however, as guaranteed by Proposition 2.

creates an option value, an effect which always dominates the first one near already explored technologies. This implies that exploitation is always dominated by slightly departing from the best explored technology.

The *value of radical innovation* is defined as the value of innovation of the unbounded unit, after any given history:

$$V^R(h_t) = V(u_\infty(h_t), z_t).$$

From Theorem 1 and Equation (6), the value of radical innovation depends on h_t only through the gap ($\rho_t = z_t - f(\bar{x}_t)$): $V^R(h_t) = \psi(\rho_t) - \rho_t$. Comparative statics for the value of radical innovation are key to understand and prove the stagnation result of Theorem 2.

PROPOSITION 3 (VALUE OF RADICAL INNOVATION) *The following properties hold after any history:*

1. *Keeping $f(\bar{x}_t)$ fixed, V^R is decreasing in ρ_t (and z_t).*
2. *If the value of radical innovation is negative at any time, radical innovation is abandoned forever after.*

To gain some intuition for Proposition 3, we observe that the frontier outcome $f(\bar{x}_t)$ entirely determines the expected outcome of any technology to the right of the frontier, and that the relative appeal of radical innovation, compared to exploitation only depends on the payoff distribution of radical innovations, relative to z . Radical innovation is appealing when it increases the agent's expected continuation payoff, which occurs only if the payoff of radical innovation exceeds z_t . A higher gap reduces the probability of this event. Keeping $f(\bar{x}_t)$ fixed, an increase in z leads to a wider gap which, in turn, depresses incentives to perform radical innovation.

Once the value of radical innovation has become negative, it will never be positive again and radical innovation will not be undertaken anymore. This insight has far-reaching consequences for the sustainability of radical innovation in the long run, and may be explained as follows: as previously argued, the value of radical innovation decreases with the size of the gap. Since the sequence $\{z_t\}$ of best explored payoffs is nondecreasing over time, and the payoff $f(\bar{x}_t)$ at the frontier is frozen whenever marginal innovation takes place, the gap $z_t - f(\bar{x}_t)$ can

only increase over time during marginal innovation, which makes radical innovation even less attractive, and creates a vicious cycle.

The second part of Proposition 3 provides a key sufficient condition for radical innovation to terminate in finite time. Theorem 2 shows that, for all sample paths, there exists a time at which this condition is satisfied.

Finally, the *value of marginal innovation* will be used to characterize innovation cycles.

DEFINITION 2 *Fix an arbitrary history h_t . The value of marginal innovation is defined by $V^M(h_t) = \max_{u \in \mathcal{P}(h_t)} V(u, z_t)$.*

It follows from these definitions that an agent prefers radical over marginal innovation if and only if the value of radical innovation dominates the value of marginal innovation.

4 Innovation Cycles

Our model gives rise to *innovation cycles*: generations alternate between radical and marginal innovations. These innovation cycles capture the Schumpeterian idea of successful innovation followed by imitation. Radical innovation generates a positive externality for all future generations by expanding the explored hull. The creation of an additional unit endows the society with new opportunities for costless marginal innovation. While marginal innovation refines knowledge about technologies in the explored hull, it is characterized by more predictable outcomes and thus lacks upside potential. As the learning value of marginal innovation goes down, radical innovation becomes attractive again, provided that the value of radical innovation remains positive. If, following a successful radical innovation, the best technology is exactly at the frontier, the gap at the next period is equal to zero. Any radical innovation is equally likely to outperform or underperform the payoff at the current frontier, and waves of radical innovations with a zero gap can occur for several periods, as long as those radical innovations are successful enough. This and other results are formalized in Proposition 4.

PROPOSITION 4 (INNOVATION CYCLES) *If, following some arbitrary history h_t , the values of marginal and radical innovations satisfy $V^M(h_t) > V^R(h_t) > 0$, then radical innovation*

takes place at some future date with positive probability. If, instead, radical innovation takes place at time t ($V^R(h_t) > V^M(h_t) > 0$) and yields outcome f , then:

1. There exists a (history-dependent) cutoff \bar{f}_t such that marginal innovation at time $t+1$ takes place in the newly created bounded unit if and only if $f > \bar{f}_t$.
2. There exists a (history-dependent) cutoff \hat{f}_t such that radical innovation at time $t+1$ is optimal if $f > \hat{f}_t$ and suboptimal if $\rho_t = 0$ and $f < \hat{f}_t$.

Thus, any wave of marginal innovation is followed by radical innovation, with positive probability, provided that the value of radical innovation is initially positive. The second part of Proposition 4 describes what follows radical innovation, depending on the outcome. If, and only if, radical innovation yields a high enough outcome, marginal innovation at the next period, if it is chosen by the next generation, takes place in the new unit. Thus, high realizations shift attention towards the newly created unit.

However, radical innovation also increases the value of further radical innovation, provided that its payoff is sufficiently high. When radical innovation is successful enough, it is always followed by radical innovation. This result is nontrivial. Indeed, a high outcome increases both the value of marginal innovation in the new unit and the value of radical innovation. Why does the latter dominate the former? Intuitively, the conditional outcome distribution on the newly created unit is roughly a straight line, with low variance.¹⁸ This means that technologies in the new unit have a lower expectation than the last technology, and a low variance. In contrast, radical innovation has the same expectation as the value of the last technology, and a variance that is independent of that level, which makes it more attractive.

Summary: Proposition 2 shows that exploitation is never optimal. Thus, any wave of marginal innovations reduces the width of the available units, as the frontier stays the same during marginal innovation. The reduction in width of these units, all else equal, also reduces the value of innovation of these units, by Theorem 1. The value of marginal innovation may then decrease to the point of triggering a new round of radical innovation, which continue until it leads to a disappointing payoff, triggering a new wave of marginal innovation to

¹⁸This result is well-known in the literature on large deviations, see e.g., Dembo and Zeitouni (1998). See also Footnote 17.

explore further the units created within the new frontier, etc. Thus, successive generations alternate between radical and marginal innovation in cycles.

5 Radical Innovation: Optimal Size and Long-Run Stagnation

Radical innovation is necessary for long-run technological growth: Without radical innovation, technology and knowledge converge to finite levels. It is therefore fundamental to determine whether radical innovation can be sustained in the long run. This section provides a negative answer, which persists even when radical innovation entails an arbitrarily positive drift, as long as the optimal policy is well-defined, as shown in Section 7.2. This section also establishes comparative statics for the optimal size of experimentation. In particular, Proposition 5 shows that the *marginal* value and the size of radical innovation decrease with the payoff of the best explored technology. by combining ideas from the experimentation literature with the spatial nature of our model.

THEOREM 2 *Radical innovation ends in finite time with probability one. After radical innovation has ended, the value of marginal innovation converges to zero almost surely, and innovations converge to a single technology.*

The end of radical innovation means that all future generations will only engage in marginal innovation, because exploitation of known technologies is always suboptimal. Thus, innovation still occurs along the equilibrium path, but the value of additional innovation keeps decreasing. The intuition is that, since the frontier stops expanding, future generations compete away the opportunities for valuable innovation. Thus, we observe the emergence of a technological standard in the limit.¹⁹

To understand why radical innovation is not sustainable in the long run, we need to analyze how the incentives for radical innovation evolve as the technological path unfolds. Those

¹⁹This result bears some resemblance to *informational cascades* analyzed in the social learning literature, especially according to the definition provided by Lee (1993) for the case of a continuum of actions. As we already mentioned in the introduction, the mechanism is different here, because the amount of information is endogenously acquired by each generation.

incentives depend on the gap: when the gap is zero, the distribution of payoffs to the right of the frontier is directly affected by the payoff z_t of the best explored technology. If the gap is strictly positive, that distribution is independent of z_t , by (1).

Let x^R denote the technology chosen when radical innovation takes place, and let $y_t^R(h_t) = x^R(h_t) - \bar{x}_t$ denote the optimal “size” of radical innovation at time t .²⁰ Let $\phi(\cdot)$ denote the density function of the standard normal distribution.²¹

PROPOSITION 5 (RADICAL INNOVATION) *Fix an arbitrary history h_t . The optimal size of radical innovation solves*

$$\frac{\delta\sigma}{2\sqrt{y}}\phi\left(\frac{\rho_t}{\sigma\sqrt{y}}\right) = c'(y) \quad (7)$$

whenever y_t^R is positive, and is equal to zero otherwise. When $\rho_t = 0$, y_t^R is the unique positive solution of (7). Finally:

1. *If $y_t^R > 0$, then y_t^R is strictly increasing in σ and δ .*
2. *If $\rho_t > 0$ and $y_t^R > 0$, then y_t^R is strictly decreasing in ρ_t .*
3. *There exists a cutoff $\tilde{\rho} > 0$ such that the value of radical innovation is negative whenever $\rho_t > \tilde{\rho}$.*

Consider, first, a period t at which the gap is equal to zero. Equation (7) captures the trade-off between marginal benefit and marginal cost of radical innovation. If a young agent chooses the best explored technology, his expected payoff is equal to $(1 + \delta)z_t$, because the agent will choose the same technology when old. If the agent chooses radical innovation, the zero drift condition implies that any positive size $y > 0$ of radical innovation has an expected payoff of exactly z_t . However, it improves the probability of a higher payoff at the next period. The marginal benefit of radical innovation (left-hand side of (7)) is large close to the frontier, making radical innovation very attractive. As the size of radical innovation increases, the volatility ($\sigma\sqrt{x - \bar{x}}$) increases at a decreasing rate, and the marginal benefit of

²⁰ There may exist several optima. In such case, the comparative statics in the proposition apply in the sense of the strong set order of lattice theory. See, e.g., Milgrom and Shannon (1994). As shown in the proof of Proposition 5, the agent’s objective function is submodular in (y, ρ) and supermodular in (y, δ) and (y, σ) .

²¹To simplify notation, the statement of Proposition 5 does not explicitly show the dependence of the optimal size of radical innovation on the history h_t .

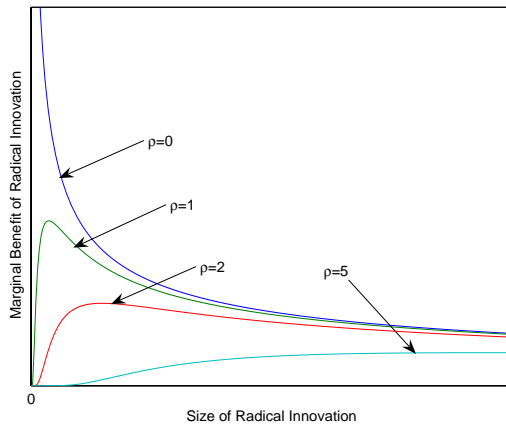


Figure 2: Marginal benefit of radical innovation.

radical innovation converges to zero. Since the marginal cost is increasing, the optimal size of radical innovation is positive and well-defined. Furthermore, an agent pursuing radical innovation will choose the same size of radical innovation when the gap is zero, regardless of the history.²²

The situation is very different when the gap is strictly positive. The marginal benefit of radical innovation now converges to zero for radical innovations close to the frontier. Close to the frontier, a radical innovation is accompanied by an inadequately low increase in volatility, and almost no impact on the expected payoff of the agent when he becomes old. The “outside” option, z_t , is thus strictly preferred to small radical innovations. In fact, the marginal benefit of radical innovation is single-peaked. It is initially pushed up by the increase in the probability of discovering an outcome above the current outside option z_t , which is given by $1 - \Phi\left(\frac{\rho_t}{\sigma\sqrt{x-\bar{x}}}\right)$, where Φ denotes the cumulative function of a standard normal distribution. When the size of radical innovation reaches $\frac{\rho^2}{\sigma^2}$, the marginal benefit starts to decrease, as the probability of obtaining an outcome greater than z_t converges $\frac{1}{2}$. Figure 2 illustrates the marginal benefit of radical innovation for different sizes of the gap.

The optimal size of radical innovation is sensitive to the volatility parameter of the underlying outcome process, because the volatility is the agent’s only way to improve his expected payoff tomorrow. A higher volatility thus has a positive effect on the incentives to innovate.

²²This follows from the Markov property of Brownian motion.

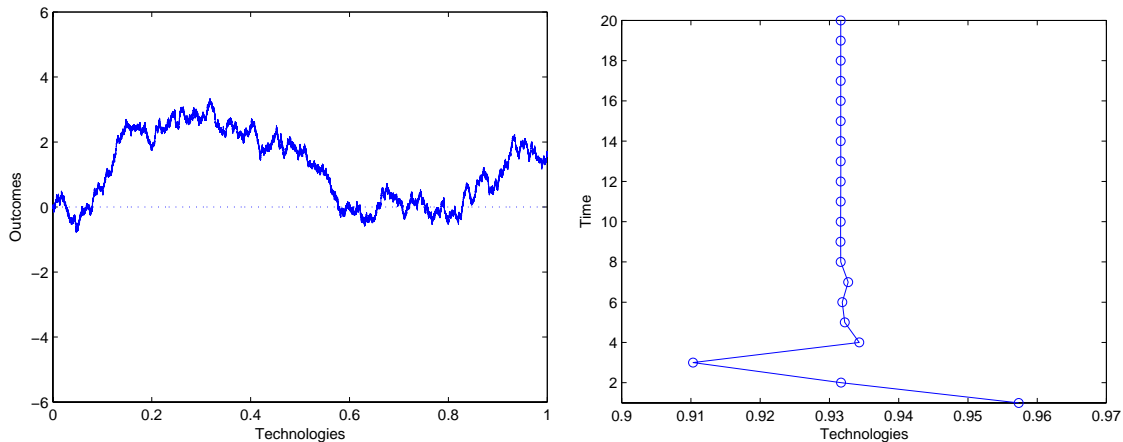
Similarly, a higher discount factor increases the incentives to innovate, given the larger weight assigned to the second-period payoff. However, an increase in the gap reduces the marginal benefit of radical innovation, because it reduces the probability of surpassing the current outside option. To maintain that probability, an agent must increase the size of radical innovation, so as to increase payoff volatility.

Proposition 5 also helps understand why radical innovation cannot be sustained in the long run: for high gaps, the optimal size of radical innovation drops, which reduces volatility and, hence, the value of radical innovation, just when volatility is most needed to make radical innovation attractive.

Proposition 5 does not alone imply that radical innovation is condemned to end in the long run, because technologies are endogenous and could a priori result in positive value of radical innovation along the equilibrium path. To prove stagnation, the proof proceeds by contradiction. If the value of radical innovation were positive at all times, the frontier would keep expanding.²³ As the frontier is moved further to the right, it will eventually hit a region where the gap exceeds $\tilde{\rho}$. Thus, radical innovation eventually ends.

Figure 3 simulates a wave of marginal innovations that follows the beginning of the stagnation phase. Figure 3(a) shows the realized path of f , while Figure 3(b) contains the results of the simulation, in which we assume that the first generation knows the value of f only at the endpoints 0 and 1. In this simulation, the first agent experiments with a technology close to 1, the endpoint with the highest outcome. The next generation moves left, surpassing earlier payoffs. The third generation also moves significantly left, receiving a lower payoff. Starting with the fourth generation, the search for a better technology begins to cluster around the technology found by the second generation, yielding the endogenous-information equivalent of an informational cascade. Society converges a suboptimal technology, which is largely path dependent: a single experiment at, say, $x = 0.3$ would completely change the dynamics of the search process, shifting innovation towards a different part of the technological space.

²³This claim is proved in the Appendix (Lemma 2), and may be coarsely understood as follows: the size of radical innovation is decreasing in the gap. We know from Proposition 5 that a high enough gap ends radical innovation forever. For radical innovation to continue, therefore, the size of the radical innovation must be bounded below.



(a) Realized path of Brownian Motion with zero drift and volatility $\sigma = 3$.

(b) Simulation

Figure 3: Dynamics of marginal innovation: $\delta = 0.95$.

6 Patents and Radical Innovation

We argued in earlier sections that radical innovation stops inefficiently early, compared to the social optimum, because each generation fails to internalize the positive information and cost spillover that radical innovation provides to future generations. We discuss here one natural way to address this inefficiency, through the use of patents. In our setting, patents consist of intergenerational transfers between consecutive generations, and are naturally introduced by exploiting the spatial nature of our model. A natural starting point is to assume that, whenever an agent undertakes radical innovation, the set of technologies that become available at no cost, due to this radical innovation, are patented to this agent.

Formally, if the time- t young generation chooses radical innovation, that generation has the right to a royalty fee if the time- $t + 1$ young generation decides to innovate marginally with any technology in the new unit with interval $\bar{u} = (\bar{x}_t, \bar{x}_t + y_t^R]$. We assume that the patent lasts for one period, which corresponds to the remaining life of the time- t generation.

A patent protection system may affect incentives for radical innovation through two channels:

Getting Royalty Fees The most obvious channel is the additional gain that an agent makes by performing radical innovation: as shown by Proposition 4, there is a positive probability that the outcome of radical innovation will be in some intermediate range at which further

radical innovation is suboptimal for the next generation, but marginal innovation takes place in the new unit. When this happens, a small royalty fee just paid by the new generation increases the ex ante incentives for radical innovation.

Perhaps counter-intuitively, even in the case of a flat royalty (let alone a more complex royalty structure), the level of that royalty affects the marginal value of radical innovation²⁴ and, therefore, the optimal size of radical innovation. Moreover, this value may be *non-monotonic* in the patent level. This may be explained as follows: a bolder radical innovation has a higher variance and, for given mean, a higher probability of reaching very high outcomes. When this happens, the next generation has a strong incentive to use the technological domain created by this radical innovation, which generates royalty fees for the old generation. A higher royalty fee may therefore increase the expected marginal value of radical innovation, spurring radical innovation and increasing the size of radical innovation. However, as the royalty fee gets arbitrarily large, the new generation prefers to forgo this opportunity and to innovate in older units, causing the non-monotonicity.

Avoiding Royalty Fees Precisely because an incoming generation has to pay a cost to the previous generation in order to innovate in a newly created unit, this reduces the value of marginal innovation for that generation, relative to radical innovation: the new generation is “pushed” towards further radical innovation.

These two incentives are substitutes: the more the new generation avoids royalty fees (say, because they are high), and the lower the patent incentives for the old generation, and vice versa. However, both of these incentives have the effect of fostering radical innovation: whether it is the perspective of getting royalty fees that increases incentives for radical innovation, or the penalty for performing marginal innovation with technologies made available by the previous generations, generations always have a higher incentive for radical innovation.

These two channels arise naturally in the spatial representation of our model, and would clearly persist under other span and duration specifications, or if the royalty depends in a more sophisticated way on past innovation (i.e., through the payoff of innovation, or the distance to explored technologies).

²⁴The term “marginal” refers here to the impact of small changes in the size of innovation on the agent’s expected benefit from innovation (its usual meaning in the common expression “marginal value”), and not to its use in the expression “marginal innovation.”

7 Extensions

7.1 Cost Externalities

The path of innovations was shown so far to affect incentives for radical innovation through: *i*) the expected value of new technologies, and *ii*) the opportunity cost of forgoing marginal innovation. In reality, radical innovation may be directly affected by past technologies, if those technologies are helpful in producing radical innovation. We model this by assuming that the cost of radical innovation is decreasing in the best available technology: the cost is given by a function $c(x - \bar{x}, \alpha z)$ for $x > \bar{x}$ with $\alpha \geq 0$ ($\alpha = 0$ corresponds to the benchmark model). We assume that, for each z , the function $c(\cdot, z)$ satisfies the earlier Assumption 1.

Let $y^R(\rho_t, \alpha z_t)$ denote the optimal size of radical innovation at a history h_t with associated gap ρ_t , and current best outcome z_t . The following result is established similarly to Proposition 5 and is stated without proof.²⁵

PROPOSITION 6 (COMPARATIVE STATICS) *Suppose that $\alpha > 0$ and that c is submodular. Then, $y^R(\rho_t, \alpha z)$ is increasing in z and in α .*

Monotonicity with respect to z of the optimal size of innovation does not necessarily imply that radical innovation itself is fostered by a higher z , even if we also assume that the cost $c(y, \alpha z)$ is decreasing in its second component. Indeed, when the gap is positive, a higher value of z reduces the cost of radical innovation, but also reduces the value of radical innovation, as shown by Proposition 3. When the gap is zero, however (following successful radical innovation), an increase in the best available technology always stimulates radical innovation. Therefore, this link between cost and the best technology should result in longer waves of radical innovation.

Monotonicity of innovation size with respect to the gap, which was established by Proposition 5, may fail in the presence of the cost externality studied here. Indeed, the reduction in the marginal benefit of radical innovation following a larger gap might be more than compensated by a decrease in the cost and marginal cost of radical innovation. Without the cost

²⁵The agent's objective function is submodular in (y, α) . When there are multiple maximizers, Proposition 6 holds in the sense of the strong set order (see footnote 20).

externality, an increase in the gap reduces the marginal benefit while leaving the marginal cost unaffected. In that case, we already know that there is a threshold for the gap above which an agent would always set the size of radical innovation to zero.

The next result shows that stagnation still occurs as long as the marginal cost of radical innovation is bounded below away from zero as the best available technology becomes arbitrarily large. We now assume that c is decreasing in its second component. Let $\bar{c}(y) = \lim_{z \rightarrow +\infty} c(y, \alpha z)$ denote the lower envelope of the cost functions $\{c(\cdot, w)\}$.²⁶

PROPOSITION 7 (STAGNATION) *Fix $\alpha > 0$, and suppose that $\bar{c}(\cdot)$ is increasing in a right neighborhood of $y = 0$. Then, radical innovation ends in finite time with probability one.*

Even if the long-run dynamics is the same with and without intergenerational cost externalities, the short-run pattern of innovation might be significantly different in the two scenarios.

7.2 Positive Drift

It was assumed, in Theorem 2, that agents are completely ignorant concerning the expected outcome of technologies to the right of the frontier, assigning them the same expected payoff as the payoff at the frontier. We now consider what happens if agents are optimistic about the underlying generating process: could a positive drift be enough to sustain radical innovation?

A positive drift complicates the analysis because it creates incentives for an old agent to innovate as well. An old agent still prefers to exploit rather than innovate marginally, but may now wish to undertake radical innovation, so as to maximize the objective

$$U^{O,R}(h_t) = \max_{x \in [\bar{x}_t, +\infty)} E_{h_t} [f(x) - c(x - \bar{x}_t)] = f(\bar{x}_t) + \kappa(x - \bar{x}_t) - c(x - \bar{x}_t) \quad (8)$$

For the problem to have an interior solution, we make the following assumption:

ASSUMPTION 2 $\lim_{y \rightarrow +\infty} c'(y) > \kappa(1 + \delta)$.

If Assumption 2 were violated, a young agent would prefer radical innovation of an arbitrarily large size, as the resulting expected lifetime benefit would exceed the immediate cost.

²⁶The function $\bar{c}(\cdot)$ is well-defined because, for each $y \geq 0$, the sequence $\{c(y, w)\}$ is decreasing in w and nonnegative.

When old agents undertake radical innovation, this changes the optimization problem of the young generation, which can learn, when old, the outcome of this innovation. The presence of another active agent clearly changes the short-run dynamics of the model. Nevertheless, radical innovation still ends with probability one.

The size of radical innovation by the old generation, whenever it takes place, is characterized by the first-order condition

$$\kappa = c'(x - \bar{x}_t) \tag{9}$$

Let $y^{O,R}$ denote the solution to (9). The optimal size of radical innovation, for an old agent, is history-independent, and strictly positive if and only if $c'(0) < \kappa$. When $c'(0) \geq \kappa$, an old agent prefers to exploit the best known technology, and the earlier analysis applies for that period. When $c'(0) < \kappa$, an old agent follows a simple cutoff rule to choose between radical and marginal innovation. Letting $\xi = \kappa y^{O,R} - c(y^{O,R}) > 0$, an old agent prefers radical innovation over exploitation if and only if

$$f(\bar{x}_t) + \kappa y^{O,R} - c(y^{O,R}) \geq z_t, \text{ or, equivalently, if } \rho_t \leq \xi. \tag{10}$$

A high gap depresses incentives to perform radical innovation even for an old generation.

THEOREM 3 *Suppose that Assumption 2 holds and that $c'(0) > 0$. Then, radical innovation stops in finite time, almost surely.*

When choosing the size of his innovation, a young agent has in mind the effect of his action today on his incentives tomorrow. This effect can be quantified in a perceived reduction of the gap from ρ_t to $\rho_t - \xi$ due to the possibility of performing radical innovation in the second period, which results in higher incentives to perform radical innovation today. However as the gap increases, ξ becomes negligible and eventually the relative benefit of radical innovation over exploitation falls short of the explicit cost of innovation. Thus the young agent will eventually opt for marginal innovation, which is still strictly better than exploitation.

8 Conclusion

This paper provides a framework for studying social learning and innovation, in which the trade-off between radical and marginal innovation determines the short-run and long-run

dynamics of technological discoveries. Unlike standard models of experimentation, technologies are correlated, but can be compartmented into conditionally independent intervals. Innovation cycles arise either through the diminishing learning value of marginal innovation, as a result of disappointing outcomes from radical innovation, or as new room for marginal innovation is created, as a result of previous radical innovation. When, as in our model, the value of radical innovation is independent from marginal innovation, radical innovation must stop in finite time because short-term profitability considerations of each generation can stall any new learning about the value of radical innovation. External intervention in the form of patents or subsidies can foster radical innovation, and the spatial setting presented here suggests natural ways of modeling such interventions.

While we focused on Brownian uncertainty with constant drift and volatility, the index characterization that we provide in Theorem 1 can be extended to other Markovian stochastic processes. Indeed, conditional independence is the key for isolating the value of a technological unit from the characteristics of other units. Such technique will not apply, however, if the payoff (or, depending on the interpretation, the productivity) of a technology is not perfectly observed after trying this technology. In that case, the payoff distribution of a given technology depends on all technologies that have been tried before, not just the two explored technologies surrounding that technology. Similarly, conditional independence is unlikely to be preserved in a spatial representation of technologies that is multidimensional, which promises interesting if challenging analysis.

We have assumed that each young generation inherited from the old generation the knowledge of all past technologies: except in our discussion on patents, we have abstracted from any loss, or cost, in the transfer of knowledge. Such transmission frictions are realistic and would be interesting to consider.²⁷

Finally, the only state variables in the economy we consider is the information transferred from one generation to another. A major extension of the analysis would allow accumulation and transfer of capital across periods, and cast the analysis into a non-stationary growth

²⁷In Acemoglu, Dahleh, Lobel, and Ozdaglar (2011) each arriving agent observes a random subset of the actions chosen by previous agents. In our setting, this would amount to each generation observing only a subset of past technologies. Niehaus (2011) explicitly considers strategic information transmission, when agents tradeoff the benefit for future generations of inheriting past knowledge and cost of conveying this knowledge.

model, bringing it closer to the standard models on technological change and growth.²⁸ A trivial extension of our model consists in each generation investing all of its capital in some technology, and f representing the total return on the technology, assuming that it is linear. This would not affect the dynamics of the analysis. A more realistic bridge between our model of innovation and growth would allow the capital stock to affect the size of innovation. The insights of the present paper about the short-term and long-term dynamics of innovation and the long-run accumulation of knowledge should prove helpful to carry out such extension.

²⁸See Acemoglu (2008) for an in-depth review of those models.

Appendix

Proof of Theorem 1

We first consider the case of a bounded unit. We have

$$\max_{x \in [0, \bar{x}]} E_{h_t} [f(x) + \delta \max\{f(x), z\}] = \max_{I \in \mathcal{I}(h_t)} \left\{ \max_{x \in I} E_u [f(x) + \delta \max\{f(x), z\}] \right\}$$

where $\mathcal{I}(h_t) = \{I' : I' \in u, \text{ for some } u \in \mathcal{P}(h_t)\}$, where $\mathcal{P}(h_t)$ is the collection of bounded units induced by h_t on $[0, \bar{x}]$, with \bar{x} denoting the current frontier. We fix a bounded unit u and suppose, without loss of generality, that $f(x_r) > f(x_l)$. Then, for any $x \in [x_l, x_r]$,

$$f(x) \sim \mathcal{N} \left(f(x_l) + \frac{f(x_r) - f(x_l)}{x_r - x_l} (x - x_l), \frac{(x - x_l)(x_r - x)}{x_r - x_l} \sigma^2 \right)$$

Letting $g(x) = \frac{x - x_l}{x_r - x_l}$,

$$f(x) - m \sim \mathcal{N} (d g(x), g(x)(1 - g(x))l\sigma^2)$$

We also define $k(x) = f(x) - m$ and $z' = z - m$ to obtain an explicit formula for the expected payoff:

$$\begin{aligned} E_u \max\{k(x), z'\} &= z' \Phi \left(\frac{z' - dg(x)}{\sigma \sqrt{g(x)(1 - g(x))l}} \right) + dg(x) \left[1 - \Phi \left(\frac{z' - dg(x)}{\sigma \sqrt{g(x)(1 - g(x))l}} \right) \right] \\ &\quad + \sigma \sqrt{g(x)(1 - g(x))l} \phi \left(\frac{z' - dg(x)}{\sigma \sqrt{g(x)(1 - g(x))l}} \right) \end{aligned}$$

where Φ and ϕ are the CDF and pdf of the standard normal distribution. This leads to

$$\begin{aligned} E_u [f(x) + \delta \max\{f(x), z\}] &= (1 + \delta)m + dg(x) + \delta E_u \max\{k(x), z'\} \\ &= (1 + \delta)m + d \left\{ g(x)(1 + \delta - \delta\Phi) + \delta \frac{z'}{d} \Phi \right. \\ &\quad \left. + \delta \sigma \sqrt{g(x)(1 - g(x))} \frac{\sqrt{l}}{d} \phi \right\} \\ &= (1 + \delta) \left\{ m + d \bar{\eta} \left(g(x), \frac{\sqrt{l}}{d}, \frac{z'}{d} \right) \right\}, \end{aligned} \tag{11}$$

where

$$\bar{\eta} \left(g(x), \frac{\sqrt{l}}{d}, \frac{z'}{d} \right) = \frac{1}{1 + \delta} \left\{ g(x)(1 + \delta - \delta\Phi) + \delta \frac{z'}{d} \Phi + \delta \sigma \sqrt{g(x)(1 - g(x))} \frac{\sqrt{l}}{d} \phi \right\}$$

The common argument of Φ and ϕ depends only on $g(x)$, \sqrt{l}/d and z'/d , as can be seen by dividing numerator and denominator by d and, hence, so does $\bar{\eta}$.

Maximizing the expected payoff over $x \in I$ yields

$$U(u; h_t) \equiv \max_{x \in I} E_u [f(x) + \delta \max\{f(x), z\}] = (1 + \delta) \left[m + d \eta \left(\frac{\sqrt{l}}{d}, \frac{z'}{d} \right) \right] \quad (12)$$

where

$$\eta \left(\frac{\sqrt{l}}{d}, \frac{z'}{d} \right) = \max_{g(x) \in [0,1]} \bar{\eta} \left(g(x), \frac{\sqrt{l}}{d}, \frac{z'}{d} \right). \quad (13)$$

Letting $\gamma(u, z) = m + d \eta \left(\frac{\sqrt{l}}{d}, \frac{z-m}{d} \right)$ yields the first part of Theorem 1.

Because

$$\begin{aligned} \frac{\partial \bar{\eta}(x, k_1, k_2)}{\partial x} &= \frac{1}{1 + \delta} \left\{ 1 + \delta - \delta \Phi \left(\frac{k_2 - x}{\sigma k_1 \sqrt{x(1-x)}} \right) + \frac{\delta \sigma k_1}{2} \frac{1 - 2x}{\sqrt{x(1-x)}} \phi \left(\frac{k_2 - x}{\sigma k_1 \sqrt{x(1-x)}} \right) \right\} \\ &\geq \frac{1}{1 + \delta} > 0 \end{aligned}$$

for any $x \leq \frac{1}{2}$, the optimum $x^*(u, z_t)$ must lie in $\left(\frac{x_r + x_l}{2}, x_r \right]$ whenever $f(x_l) < f(x_r)$.

For this and other comparative statics results, we will apply differential techniques to the relevant value functions. These techniques can be applied because the relevant functions are always left and right differentiable, a result which follows from Corollary 4 of Milgrom and Segal (2002).²⁹ All the comparative statics arguments to follow can be applied to the right derivative (or, in case of a decrease of the parameter, to the left derivative). For expositional simplicity, we drop reference to the side of the derivative. The Envelope Theorem implies, for any bounded unit u , that³⁰

$$\frac{\partial \gamma(u, z)}{\partial l} = \frac{\delta}{1 + \delta} \frac{\sigma}{2} \sqrt{\frac{x^*(1-x^*)}{l}} > 0 \quad (14)$$

and

$$\frac{\partial \gamma(u, z)}{\partial d} = \frac{x^*}{1 + \delta} + \frac{\delta}{1 + \delta} x^* \left[1 - \Phi \left(\frac{z' - dx}{\sigma \sqrt{x^*(1-x^*)} l} \right) \right] > 0. \quad (15)$$

²⁹The value function may fail to be differentiable at parameter values for which there exist multiple maximizers. At such values, however, the value function is still left and right differentiable, with the left (right) derivative being evaluated at the maximizer that minimizes (maximizes) the derivative of the objective with respect to the parameter. See Milgrom and Segal (2002).

³⁰An argument similar to the proof of Proposition 2, which holds independently of the present comparative statics, shows that any optimum x^* is always in the interior of a unit.

For the unbounded unit u_∞ , the expected utility from a technology $x > \bar{x}$ is

$$\begin{aligned} E_{h_t}[f(x) - c(x - \bar{x}) + \delta \max\{f(x), z\}] &= (1 + \delta)f(\bar{x}) - c(x - \bar{x}) \\ &\quad + \delta E_{h_t} \max\{f(x) - f(\bar{x}), z - f(\bar{x})\} \\ &= (1 + \delta)f(\bar{x}) - c(x - \bar{x}) + \delta E_{h_t} \max\{k(x), \rho\}, \end{aligned}$$

where $k(x) = f(x) - f(\bar{x})$. Recall that $k(x) \sim \mathcal{N}(0, \sigma^2(x - \bar{x}))$. Therefore,

$$E_{h_t} \max\{k(x), \rho\} = \sigma\sqrt{x - \bar{x}}\phi\left(\frac{\rho}{\sigma\sqrt{x - \bar{x}}}\right) + \rho\Phi\left(\frac{\rho}{\sigma\sqrt{x - \bar{x}}}\right). \quad (16)$$

Taking the supremum over $x \geq \bar{x}$, we obtain

$$\begin{aligned} &\sup_{x \geq \bar{x}} \left\{ (1 + \delta)f(\bar{x}) - c(x - \bar{x}) + \delta\sigma\sqrt{x - \bar{x}}\phi\left(\frac{\rho}{\sigma\sqrt{x - \bar{x}}}\right) \right. \\ &\quad \left. + \delta\rho\Phi\left(\frac{\rho}{\sigma\sqrt{x - \bar{x}}}\right) \right\} \\ &= (1 + \delta) \left[f(\bar{x}) + \sup_{y \geq 0} \frac{1}{1 + \delta} \left\{ -c(y) + \delta\sigma\sqrt{y}\phi\left(\frac{\rho}{\sigma\sqrt{y}}\right) + \delta\rho\Phi\left(\frac{\rho}{\sigma\sqrt{y}}\right) \right\} \right]. \end{aligned}$$

Defining

$$\psi(\rho) = \frac{1}{1 + \delta} \sup_{y \geq 0} \left\{ -c(y) + \delta\sigma\sqrt{y}\phi\left(\frac{\rho}{\sigma\sqrt{y}}\right) + \delta\rho\Phi\left(\frac{\rho}{\sigma\sqrt{y}}\right) \right\} \quad (17)$$

and $\gamma(u, z) = f(\bar{x}) + \psi(\rho)$ yields the formula claimed by Theorem 1. The index of each unit corresponds to a normalization of the agent's value function when the domain of choice is restricted to that particular unit. Therefore, the agent optimally chooses a technology within the unit with the highest index. \blacksquare

Proof of Proposition 1

For any bounded unit u with $M \neq z$ we have³¹

$$\frac{d\eta(k_1, k_2(z))}{dz} = \frac{1}{d} \frac{\partial\eta(k_1, k_2)}{\partial k_2} = \frac{\delta}{(1 + \delta)d} \Phi\left(\frac{k_2 - x^*}{\sigma k_1 \sqrt{x^*(1 - x^*)}}\right)$$

where $k_1 \equiv \frac{\sqrt{l}}{d}$, $k_2 \equiv \frac{z - m}{d}$, and $x^* \equiv x^*(u, z)$ is an optimal technology within u . Thus, $\frac{\partial V(u, z)}{\partial z} = \frac{\delta}{1 + \delta} \Phi\left(\frac{z - m - dx^*}{\sigma\sqrt{x^*(1 - x^*)}l}\right) - 1 < 0$.

³¹Again, we omit dependence on side-derivatives. See Footnote 29 and the discussion surrounding it.

If, instead, u is such that $M = z$, then the index of the unit is given by $\gamma(u, z) = m + (z - m)\eta\left(\frac{\sqrt{l}}{z-m}, 1\right)$ which gives $\frac{\partial V(u, z)}{\partial z} = \eta\left(\frac{\sqrt{l}}{z-m}, 1\right) - \frac{\sqrt{l}}{z-m} \frac{\partial \eta}{\partial k_1}\left(\frac{\sqrt{l}}{z-m}, 1\right) - 1$, which can be shown to be strictly negative from (22) in the proof of Proposition 4. Then, $\frac{\partial V^M(h_t)}{\partial z} \leq \max_{u \in \mathcal{P}(h_t)} \frac{\partial V(u, z)}{\partial z} < 0$.

Exploitation of the best explored technology is always feasible for a young agent, and yields a payoff of $(1 + \delta)z_t$. Consider a unit u such that $V(u, z_t) < 0$. Then,

$$\gamma(u, z_t) < z_t \iff \frac{U(u; h_t)}{1 + \delta} < z_t \iff U(u; h_t) < (1 + \delta)z_t$$

where the second equivalence follows from (12). Therefore, choosing the best current technology dominates choosing any technology in u . Since z_t is nondecreasing in t , once the value of innovation of a unit is negative, it remains negative. ■

To simplify the analysis, we define the indexes for marginal and radical innovation by

$$\gamma^M(h_t) = \max_{u \in \mathcal{P}(h_t)} \gamma(u, z) \quad \text{and} \quad \gamma^R(h_t) = \gamma(u_\infty, z) \quad (18)$$

for any history h_t , where $\mathcal{P}(h_t)$ is the collection of bounded units induced by a history h_t .

Proof of Proposition 2

We prove only the first part, the second part being shown in the main text. Suppose that $M = z$. Differentiating $\bar{\eta}(x, k_1, 1)$ with respect to x yields

$$(1 + \delta) \frac{\partial \bar{\eta}(x, k_1, 1)}{\partial x} = \left[1 + \delta - \delta \Phi \left(\frac{1}{\sigma k_1} \sqrt{\frac{1-x}{x}} \right) \right] + \frac{\delta \sigma k_1}{2} \frac{1-2x}{\sqrt{x(1-x)}} \phi \left(\frac{1}{\sigma k_1} \sqrt{\frac{1-x}{x}} \right)$$

which tends to $-\infty$ as x goes to 1. Also, $\bar{\eta}(0, k_1, 1) \equiv \lim_{x \rightarrow 0} \bar{\eta}(x, k_1, 1) = \frac{\delta}{1+\delta} < 1 = \lim_{x \rightarrow 1} \bar{\eta}(x, k_1, 1) \equiv \bar{\eta}(1, k_1, 1)$. We have thus shown that any solution to (13) lie in $(0, 1)$ whenever $k_1 > 0$ and $k_2 = 1$. Consider any history h_t and unit u with z_t as one its endpoint payoffs. Then,

$$\gamma(u, z_t) = m + d\eta \left(\frac{\sqrt{l}}{d}, \frac{z_t - m}{d} \right) = m + d\eta \left(\frac{\sqrt{l}}{d}, 1 \right) > m + d\bar{\eta} \left(1, \frac{\sqrt{l}}{d}, 1 \right) = m + d = z$$

which proves that there is always at least one bounded unit with strictly positive value of innovation. Since $V^M(h_t) \geq V(u, z_t)$, the result follows. ■

Proof of Proposition 3

Fix $f = f(\bar{x})$ and consider the effect of an increase in $z > f$ on the value of radical innovation $V^R(h_t)$. We suppose once again that the value function (17) is differentiable. If there exists an interior solution $y^* > 0$ to (17), then $\frac{\partial V^R(h_t)}{\partial z} = \frac{d\psi(z-f)}{d\rho} - 1 = \frac{\delta}{1+\delta} \Phi\left(\frac{z-f}{\sigma\sqrt{y^*}}\right) - 1 < 0$. If there is no interior solution, then $\psi(\rho) = \frac{\delta}{1+\delta}(z-f)$ which gives $\frac{\partial V^R(h_t)}{\partial z} = \frac{\delta}{1+\delta} - 1 < 0$. When $\rho = 0$, $V^R(h_t) = \psi(0)$ and $\frac{\partial V^R(h_t)}{\partial z} = 0$. If $V^R(h_t) < 0$, the previous comparative statics, monotonicity of the sequence $\{z_t\}$, and the fact that $f(\bar{x}_t)$ is unaffected by marginal innovation, imply that radical innovation will never be undertaken for any $t' > t$, as it is always dominated by exploitation of the highest known payoff. ■

Proof of Proposition 4

Let h_t denote any history such that $V^M(h_t) > V^R(h_t) > 0$. By continuity, there exists $\varepsilon > 0$ such that the value of radical innovation remains positive if the current gap is increased to $\rho_t + \varepsilon$. Let $B(\varepsilon)$ denote the set of paths of f on $[0, \bar{x}_t]$ which are compatible with h_t and are bounded above by $z_t + \varepsilon$. $B(\varepsilon)$ occurs with positive probability. By construction, the sequence $\{z_{t'}\}_{t' > t}$ stays below $z_t + \varepsilon$ for any path in $B(\varepsilon)$, which implies that the value of marginal innovation converges to zero over time, by Theorem 2. Thus, the value of marginal innovation must fall below the value of radical innovation, which is uniformly bounded away from zero for any path in B .

1.) Let f denote the outcome of radical innovation following history h_t , and let $u^{new}(f)$ be the new bounded unit generated following radical innovation. The new index for marginal innovation is

$$\gamma^M(h_{t+1}) = \max \left\{ \underbrace{\max_{u \in \mathcal{P}(h_t)} \gamma(u, \max\{z_t, f\})}_{\bar{\gamma}(\max\{z_t, f\})}, \gamma(u^{new}(f), \max\{z_t, f\}) \right\}$$

We first consider how $\bar{\gamma}(\max\{z_t, f\})$ varies with f . If $f \leq z_t$, $\bar{\gamma}(\cdot)$ is unaffected by f . If $f > z_t$, consider any unit $u \in \mathcal{P}(h_t)$. Then,³²

$$\frac{\partial \gamma(u, f)}{\partial f} = \frac{\partial}{\partial f} \left[m + d\eta \left(\frac{\sqrt{l}}{d}, \frac{f-m}{d} \right) \right] = \frac{\partial \eta}{\partial k_2} \left(\frac{\sqrt{l}}{d}, \frac{f-m}{d} \right) \in \left(0, \frac{\delta}{1+\delta} \right) \quad (19)$$

where $x^* \in \arg \max_{x \in [0,1]} \bar{\eta} \left(x, \frac{\sqrt{l}}{d}, \frac{f-m}{d} \right)$ is an optimal technology for the unit u .

³²Again dropping dependence on side-derivatives, see Footnote 29.

Next, we consider the index of the new unit u^{new} . To this end, define f_a as the outcome associated with the left endpoint of unit u^{new} . If $f_a < f \leq z_t$, then

$$\begin{aligned} \frac{\partial \gamma(u^{new}(f), z_t)}{\partial f} &= \frac{\partial}{\partial f} \left[f_a + (f - f_a) \eta \left(\frac{\sqrt{l}}{f - f_a}, \frac{z_t - f_a}{f - f_a} \right) \right] \\ &= \frac{\hat{x}^*}{1 + \delta} \left[1 + \delta - \delta \Phi \left(\frac{k_2 - \hat{x}^*}{\sigma k_1 \sqrt{\hat{x}^*(1 - \hat{x}^*)}} \right) \right] > 0 \end{aligned} \quad (20)$$

where \hat{x}^* is an optimal technology inside the new unit. If, instead, $f < \min\{z_t, f_a\}$, then

$$\begin{aligned} \frac{\partial \gamma(u^{new}(f), z_t)}{\partial f} &= \frac{\partial}{\partial f} \left[f + (f_a - f) \eta \left(\frac{\sqrt{l}}{f_a - f}, \frac{z_t - f}{f_a - f} \right) \right] \\ &= (1 - \hat{x}^*) \left[1 - \frac{\delta}{1 + \delta} \Phi \left(\frac{k_2 - \hat{x}^*}{\sigma k_1 \sqrt{\hat{x}^*(1 - \hat{x}^*)}} \right) \right] \in \left(0, \frac{1}{1 + \delta} \right) \end{aligned} \quad (21)$$

The upper bound follows from the fact that $\hat{x}^* \in (\frac{1}{2}, 1)$, then

$$(1 - \hat{x}^*) \left[1 - \frac{\delta}{1 + \delta} \Phi \left(\frac{k_2 - \hat{x}^*}{\sigma k_1 \sqrt{\hat{x}^*(1 - \hat{x}^*)}} \right) \right] \leq \frac{1}{2} \left[1 - \frac{\delta}{1 + \delta} \frac{1}{2} \right] = \frac{2 + \delta}{4(1 + \delta)} < \frac{1}{1 + \delta}$$

for any $\delta \in [0, 1]$. Finally, suppose that $f > z_t$ ($\geq f_a$), then we have

$$\frac{\partial \gamma(u^{new}(f), f)}{\partial f} = \frac{1}{1 + \delta} \left[\hat{x}^*(1 + \delta) + \delta(1 - \hat{x}^*) \Phi \left(\frac{1}{\sigma k_1} \sqrt{\frac{1 - \hat{x}^*}{\hat{x}^*}} \right) \right] \in \left(\frac{\delta}{1 + \delta}, 1 \right) \quad (22)$$

The upper bound follows from $\hat{x}^* < 1$, since $k_2 = 1$ in this case, while the lower bound follows from $\hat{x}^* \geq \frac{1}{2}$, by Theorem 1.

Below z_t , $\bar{\gamma}$ is flat, while the index of the new unit strictly increases in f , by (20) and (21). Above z_t , $\bar{\gamma}$ grows at most by $\frac{\delta}{1 + \delta}$ following an increase in f , while the derivative of the index of the new unit is given by (22). If $\bar{\gamma}(z_t) > \gamma(u^{new}(z_t), z_t)$, the index of the new unit is strictly below $\bar{\gamma}(z_t)$ for any $f < z_t$. Since above z_t , the lowest slope of the index of the new unit is strictly higher than the largest slope of $\bar{\gamma}(\cdot)$, the two indexes intersect exactly once. If instead $\bar{\gamma}(z_t) < \gamma(u^{new}(z_t), z_t)$, the two indexes necessarily cross only once at some $f < z_t$, but never above z_t . We ignore the case $\bar{\gamma}(z_t) = \gamma(u^{new}(z_t), z_t)$, which occurs with zero probability.

2.) Let f denote the new frontier outcome. With a slight abuse of notation, the index of radical innovation at time $t + 1$ is given by $\gamma_{t+1}^R(f) = f + \psi(\rho_{t+1}) = f + \psi(\max\{z_t - f, 0\})$ as a function of the realized outcome f . For any $f \geq z_t$, $\gamma_{t+1}^R(f) = f + \psi(0)$ and $\frac{\partial \gamma_{t+1}^R}{\partial f} = 1$.

Suppose instead that $f < z_t$. From Proposition 5, there exists a threshold $\tilde{\rho}$ above which the optimal size of radical innovation is 0. Define \tilde{f} by $z_t - \tilde{f} = \tilde{\rho}$. Then, for any $f \leq \tilde{f}$, $\gamma_{t+1}^R(f) = \frac{f + \delta z_t}{1 + \delta}$, because the time- $t + 1$ generation would prefer the frontier technology to any technology to its right. Thus, in this case $\frac{\partial \gamma_{t+1}^R}{\partial f} = \frac{1}{1 + \delta}$.

For $f \in (\tilde{f}, z_t)$, the maximizer of (17) may be interior. In that case, $\frac{\partial \gamma_{t+1}^R}{\partial f} = 1 - \frac{\delta}{1 + \delta} \Phi \left(\frac{z_t - f}{\sigma \sqrt{y^*}} \right) \in \left(\frac{1}{1 + \delta}, 1 \right)$, where y^* is the optimal size of radical innovation following history h_{t+1} . Thus, the derivative is always at least $\frac{1}{1 + \delta}$.

We need to consider two cases.

Case 1: $\rho_t = 0$. Let \bar{f} be the cutoff derived in Proposition 4.1), which we just proved. If $\bar{f} \geq z_t$, the index of marginal innovation is flat for any $f < z_t$, and

$$\frac{\partial \gamma_{t+1}^M}{\partial f} \in \begin{cases} \left(0, \frac{\delta}{1 + \delta}\right) \text{ from (19)} & \text{if } z_t \leq f < \bar{f} \\ \left(\frac{\delta}{1 + \delta}, 1\right) \text{ from (22)} & \text{if } f \geq \bar{f} \end{cases}$$

which follows from our previous analysis. Since the slope of the index of radical innovation is always at least $\frac{1}{1 + \delta} > \frac{\delta}{1 + \delta}$ below \bar{f} and 1 above \bar{f} , it follows that γ_{t+1}^R and γ_{t+1}^M as functions of f cross exactly once. Let $\hat{f} = \hat{f}(h_t)$ denote such intersection.

Similarly, if $\bar{f} < z_t$, the index for radical innovation is unchanged, but

$$\frac{\partial \gamma_{t+1}^M}{\partial f} \begin{cases} = 0 & \text{if } f < \bar{f} \\ \in \left(0, \frac{1}{1 + \delta}\right) \text{ from (21)} & \text{if } \bar{f} \leq f < z_t \\ \in \left(\frac{\delta}{1 + \delta}, 1\right) \text{ from (22)} & \text{if } f \geq z_t \end{cases}$$

A direct comparison of the slopes of the indexes shows that there exists exactly one intersection.

Case 2: $\rho_t > 0$. If $\bar{f} \geq z_t$, the analysis is the same as for Case 1. Thus, there exists a unique intersection \hat{f} .

If $\bar{f} < z_t$, the slope of the index of marginal innovation over the range $[\bar{f}, z_t)$ is given by (20), which cannot be compared with the slope of γ_{t+1}^R in an unambiguous way. Thus, we cannot exclude the possibility of multiple intersections between the two indexes.

Proof of Proposition 5

If $\rho = 0$, the expected utility from radical innovation of size $y > 0$ is equal to

$$E_{h_t} [f(\bar{x}_t + y) - c(y) + \delta \max\{f(\bar{x}_t + y), z_t\}] = (1 + \delta)f(\bar{x}_t) - c(y) + \delta\sigma\sqrt{y}\phi(0),$$

from (16). The first-order condition yields Equation (7) with $\rho = 0$. The right-hand side of (7) is increasing in y , while the left-hand side is strictly decreasing. The left-hand side is also unbounded around 0, and converges to 0 as $y \rightarrow +\infty$. Therefore, there always exists a solution to Equation (7) when $\rho = 0$, and it is unique. The second-order condition

$$-\frac{\delta\sigma\phi(0)}{4y^{3/2}} - c''(y) < 0$$

is satisfied, guaranteeing that the first-order condition characterizes maxima.

If $\rho > 0$, the expected utility from radical innovation of size $y > 0$ is

$$(1 + \delta)f(\bar{x}) - c(y) + \delta\sigma\sqrt{y}\phi\left(\frac{\rho}{\sigma\sqrt{y}}\right) + \delta\rho\Phi\left(\frac{\rho}{\sigma\sqrt{y}}\right)$$

The first-order condition is again given by (7).

1.) and 2.) Differentiating (7) with respect to ρ , we obtain $-\frac{\delta\rho}{2\sigma y^{3/2}}\phi\left(\frac{\rho}{\sigma\sqrt{y}}\right)$, which is strictly negative. Thus, the objective function in (17) is submodular in (y, ρ) , and the result follows from the Strict Monotonicity Theorem 1 of Edlin and Shannon (1998). Similarly, the objective function is also supermodular in (y, δ) and (y, σ) , which completes the proof.

3.) We start the analysis with some preliminary results. First, let $A(y, \rho) \equiv \frac{\delta\sigma}{2\sqrt{y}}\phi\left(\frac{\rho}{\sigma\sqrt{y}}\right)$, which denotes the marginal benefit of radical innovation given a gap ρ and a size y of radical innovation.³³

LEMMA 1 *For any $\rho > 0$,*

1. $\lim_{y \rightarrow +\infty} A(y, \rho) = \lim_{\rho \rightarrow +\infty} A(y, \rho) = \lim_{y \rightarrow 0} A(y, \rho) = \lim_{y \rightarrow 0} \frac{\partial A(y, \rho)}{\partial y} = 0;$
2. $\frac{\partial A(y, \rho)}{\partial \rho} < 0.$
3. $A(y, \rho)$ is strictly quasi-concave in y , and maximized at $\frac{\rho^2}{\sigma^2}.$

Proof. The first two limits in 1) and the sign of the derivative in 2) directly follow from the properties of $\phi(\cdot)$. The limit $\lim_{y \rightarrow 0} A(y, \rho)$ is computed using the fact that $\lim_{y \rightarrow 0} \frac{1}{\sqrt{y}}e^{-\frac{1}{y}} =$

³³ $A(y, \rho)$ corresponds to the left-hand side of equation (7).

$\lim_{z \rightarrow \infty} \frac{z^{1/2}}{e^z} = 0$. One shows similarly that $\lim_{y \rightarrow 0} \frac{\partial A(y, \rho)}{\partial y} = 0$. Strict quasiconcavity of A in y comes from the fact that

$$\frac{\partial A(y, \rho)}{\partial y} = \frac{\delta \sigma}{4y^{3/2}} \left[\frac{\rho^2}{\sigma^2 y} - 1 \right] \phi \left(\frac{\rho}{\sigma \sqrt{y}} \right),$$

which is nonnegative below $\frac{\rho^2}{\sigma^2}$ and negative above. This also shows that $\frac{\rho^2}{\sigma^2}$ maximizes $A(\cdot, \rho)$. ■

We now show the existence of a threshold $\tilde{\rho} > 0$ above which radical innovation is strictly suboptimal. We start by showing that $y^* = 0$ is the unique maximizer of (17), whenever ρ is high enough. If $c'(0) > 0$, Lemma 1 implies that $\lim_{y \rightarrow 0} A(y, \rho) = 0 < c'(0)$ for any $\rho > 0$. Also by Lemma 1, $\frac{\partial A(y, \rho)}{\partial \rho} < 0$, $\lim_{\rho \rightarrow +\infty} A(y, \rho) = 0$, and $A(y, \rho) \leq \frac{\rho^2}{\sigma^2}$. This, combined with the properties of c , implies the existence of a large enough threshold $\tilde{\rho}$ such that $c'(y) > A(y, \rho)$ for all $\rho > \tilde{\rho}$ and $y > 0$. A similar argument applies if $c'(0) = 0$ and $c''(0) > 0$, because $\lim_{y \rightarrow 0} \frac{\partial A(y, \rho)}{\partial y} = 0$, from Lemma 1. Substituting $y^* = 0$ into (17) yields $\psi(\rho) = \frac{\delta \rho}{1+\delta}$. Therefore, the index of radical innovation is equal to $\gamma^R = f(\bar{x}) + \frac{\delta \rho}{1+\delta} = \frac{f(\bar{x}) + \delta z}{1+\delta}$, which is strictly less than z , for $\rho > 0$.

Proof of Proposition 7

$\bar{c}(\cdot)$ is increasing in a right neighborhood of $y = 0$. Assumption 1, which applies to all functions $c(\cdot, z)$, then guarantees that $\bar{c}(\cdot)$ is increasing everywhere. Replicating the proof of Proposition 5, one can show the existence of a threshold $\tilde{\rho} > 0$ above which the marginal benefit of radical innovation is strictly less than the marginal cost, at any $y > 0$, for any $z > 0$. Thus, $y^R(\rho, \alpha z) = 0$ for any $\rho > \tilde{\rho}$, regardless of the absolute level of z .

For fixed z , $y^R(0, \alpha z)$ is still an upper bound on the size of radical innovation for any $\rho > 0$, because the incentives to perform radical innovation are maximal with a zero gap. Moreover, $y^R(0, \alpha z)$ increases in z , because the cost function is submodular in $(y, \alpha z)$. However, the marginal cost is bounded away from zero by the properties of the lower envelope $\bar{c}(\cdot)$, which implies that there exists $0 < \Lambda < +\infty$ such that $\lim_{z \rightarrow +\infty} y^R(0, \alpha z) < \Lambda$. Finally, we can now repeat the same proof as in Step 1 of Theorem 2.

9 Proof of Theorem 2

Step 1: We first show that radical innovation ends in finite time almost surely. From Proposition 5, there exists a threshold $\tilde{\rho} > 0$ such that radical innovation ends after any history such that $\rho_t \geq \tilde{\rho}$. Therefore, it suffices to show that this threshold is reached almost surely.

Consider an innovation path along which radical innovation happens infinitely often, and let $\{\varphi(t)\}$ denote the sequence of times at which radical innovation occurs. In particular, $y_{\varphi(t)}^R = x_{\varphi(t)}^R - \bar{x}_{\varphi(t)} > 0$ for all t .

LEMMA 2 $\{\bar{x}_{\varphi(t)}\}$ is unbounded a.s.

Proof. Suppose that $\bar{x}_{\varphi(t)}$, which is increasing in t , converges to some finite limit \tilde{x} . This implies that $y_{\varphi(t)}^R$ converges to 0. From Proposition 5, $y_{\varphi(t)}^R$ is decreasing in the gap. Therefore, there must exist a subsequence $\{\psi(t)\}$ of $\{\varphi(t)\}$ such that $\{\rho_{\psi(t)} = z_{\psi(t)} - f(\bar{x}_{\psi(t)})\}$ is increasing. Since that sequence is bounded above by $\tilde{\rho}$, it must converge to some strictly positive limit \tilde{g} , and $z_{\psi(t)}$ converges to the limit $\tilde{g} + f(\tilde{x})$. For sufficiently high t , the expected outcome from radical innovation is approximately equal to $f(\tilde{x}) + \delta\tilde{z}$, while the outcome from exploitation is approximately $(1 + \delta)\tilde{z}$. Since $\tilde{z} > f(\tilde{x})$ by construction, an agent will eventually strictly prefer to exploit the technology yielding \tilde{z} over the radical innovation corresponding to times in $\{\varphi(t)\}$, a contradiction. ■

From Proposition 5, $\bar{y} = y^R(0)$ is an upper bound on the optimal size of radical innovation for any size of the gap. Therefore, $|\bar{x}_{\varphi(t+1)} - \bar{x}_{\varphi(t)}| \leq \bar{y}$. For any $\zeta > 0$ and path f , let

$$A_{\zeta}(-\tilde{\rho}) = \sup \left\{ x' - x : \max_{x'' \in [x, x']} \{f(x'')\} < -\tilde{\rho}, \text{ and } x < x' < \zeta \right\}$$

LEMMA 3 $A_{\zeta}(-\tilde{\rho}) > \bar{y}$ almost surely as $\zeta \rightarrow +\infty$.

Proof. By the recurrence property of Brownian motion, there exists a.s. an $\tilde{x} > 0$ such that $f(\tilde{x}) < -\tilde{\rho}$ and, hence, some $\zeta > 0$ such that $A_{\zeta}(-\tilde{\rho}) > 0$. The result then follows from the scaling property of Brownian motion. ■

Lemma 3 means that if radical innovation goes far enough, with each leap size bounded above by \bar{y} , the frontier is bound to “fall” into a region where its payoff is less than $-\tilde{\rho}$. Because z is always nonnegative, the gap $z - f(\bar{x})$ after such history will exceed $\tilde{\rho}$, prompting

radical innovation to stop. Combined with Lemma 2, this guarantees that radical innovation must stop in finite time, almost surely.

Step 2: Belief convergence.

After radical innovation stops, all innovation takes place in a compact domain. The payoff distribution over that domain is characterized by finitely many Brownian bridges, whose endpoints correspond to previously explored technologies. We now establish that the beliefs resulting from the subsequent innovation converge to a well-defined limit. In later steps, we will show that innovations also converge to a single technology. Let $K = [0, \bar{x}]$ denote the domain of innovation after radical innovation has stopped, and μ_0 denote the distribution of f on K , given the history leading up to the end of radical innovation. For notational simplicity, we will reset to 0 the time at which radical innovation has stopped.

Let Θ denote the space of continuous functions on K starting at 0. At any time t the belief μ_t is a probability distribution over Θ : $\mu_t \in \Delta(\Theta)$. Some arguments that we need to use hold only for compact spaces. Because of this, we will sometimes need to replace the paths f by some bounded counterpart. For $\Lambda > 1$, we will consider any transformation $G(\cdot, \Lambda)$ of \mathbb{R} such that i) $G(\cdot, \Lambda)$ is continuous and strictly increasing, ii) $G(x, \Lambda) = x$ for all x such that $|x| < \Lambda - 1$, and iii) $\lim_{x \rightarrow -\infty} G(x, \Lambda) = -\Lambda$ and $\lim_{x \rightarrow +\infty} G(x, \Lambda) = +\Lambda$. Such function is easily built, and is bounded by Λ .

For any Brownian path f , the transformed path $g^\Lambda : x \mapsto G(f(x), \Lambda)$ is continuous and bounded by Λ , and is homeomorphic to f . In particular, g^Λ and f are observationally equivalent. Any belief μ about f has a corresponding belief μ^Λ about g^Λ and vice versa. Let $\Theta(\Lambda)$ denote the subset of Θ whose elements are bounded in absolute value by Λ , and $\Delta(\Lambda)$ denote the set of distributions over $\Theta(\Lambda)$

Given a sequence $\{x_t\}_{t \geq 0}$ of technology choices, let $\{\mu_t^\Lambda\}$ denote the sequence of beliefs in $\Delta(\Lambda)$ about the underlying transformed path g^Λ , obtained through Bayesian updating. It is well-known that this sequence is a martingale and converges to some limit μ^Λ . This result follows from the Martingale Convergence Theorem, and is proved similarly to Theorem 4 in Easley and Kiefer (1988). Translating this result in terms of f , this shows that the sequence $\{\mu_t\}$ of beliefs about the path f also converge to some limit μ .

For any history h leading to the belief μ , let $Z^\Lambda(\mu) = \sup\{g^\Lambda(x_t) : x_t \text{ contained in } h\}$. As

is easily checked, $Z^\Lambda(\mu)$ is independent of the particular history leading up to the limiting belief μ , and continuous in μ . We can similarly let $Z(\mu) = \max\{f(x_t)\} = G^{-1}(Z^\Lambda(\mu), \Lambda)$ ³⁴ where, for each Λ , $G^{-1}(\cdot, \Lambda)$ denotes the inverse of $G(\cdot, \Lambda)$

Step 3: Technology Convergence and Vanishing Value of Marginal Innovation.

The next step is to characterize the limit to which technologies converge. For any $(z_1, z_2) \in \mathbb{R} \times \mathbb{R}_+$, let $r(z_1, z_2) = z_1 + \delta \max\{z_1, z_2\}$ denote the payoff of an agent if the outcome of his chosen technology when young is z_1 and the best explored payoff until then was z_2 . Given a technology x , payoff z , and belief μ , let

$$u(x, \mu, z) = \int_{\Theta} r(f(x), z) d\mu(f).$$

and

$$u^\Lambda(x, \mu, z) = \int_{\Theta} r(G(f(x), \Lambda), z) d\mu(f).$$

Using the distribution μ^Λ implied on $\Theta(\Lambda)$ by μ , we have

$$u^\Lambda(x, \mu, z) = v(x, \mu^\Lambda, z),$$

where we define v , for any $\tilde{\mu} \in \Delta(\Lambda)$, by

$$v(x, \tilde{\mu}, z) = \int_{\Theta(\Lambda)} r(g(x), z) d\tilde{\mu}(g).$$

We will use the following lemma, which is proved at the end of this Appendix (Section 9.1):

LEMMA 4 $v(x, \mu, z)$ is continuous over $K \times \Delta(\Lambda) \times [-\Lambda, \Lambda]$.

Given a belief μ with corresponding maximum explored payoff z , a young agent solves the maximization problem.

$$U(\mu, z) = \max_{x \in K} u(x, \mu, z)$$

The equilibrium technological path, denoted $\{x_t^*\}$ is such that, for each t , x_t^* maximizes $u(x, \mu_t, z_t)$. We now derive properties for the long-run technologies arising in equilibrium. Given a sequence of technologies $\{x_t\}_{t=0}^\infty$, let $\mathcal{M}(\{x_t\})$ be the set of its limit points.

PROPOSITION 8 For any history h , limiting belief μ , and $x \in \mathcal{M}(\{x_t^*\})$, $x \in \operatorname{argmax}_{x' \in K} u(x', \mu, Z(\mu))$ and $f(x) = Z(\mu)$.

³⁴This maximum is also well defined, because f is continuous on the compact domain K .

Proof. Let $\{x_{t_k}^*\}$ denote a subsequence converging to x . By construction,

$$u(x_{t_k}^*, \mu_{t_k}, Z(\mu_{t_k})) \geq u(x', \mu_{t_k}, Z(\mu_{t_k})) \quad (23)$$

for any $x' \in K$. Because Lemma 4 applies only to bounded payoffs, we cannot directly take the limit in the previous inequality. Instead we will approximate it by its equivalent when the payoffs are bounded by Λ for Λ arbitrarily large. Let $\Omega(\Lambda) = \{f \in \Theta : \max_{x \in K} |f(x)| > \Lambda - 1\}$. We have for any $x, \hat{\mu}, z$

$$|u(x, \hat{\mu}, z) - u^\Lambda(x, \hat{\mu}, z)| \leq \int_{\Theta} |r(f(x), z) - r(G(f(x), \Lambda), z)| d\hat{\mu}(f) \leq 2 \int_{\Omega(\Lambda)} |r(f(x), z)| d\hat{\mu}(f). \quad (24)$$

We now show that the right-hand side converges to zero as Λ goes to infinity, uniformly on the domain $K \times \cup_t \{\mu_t\} \times [0, Z(\mu)]$. For all $x \in K$ and $z \in [0, Z(\mu)]$, $|r(f(x), z)| \leq (1 + \delta)(Z(\mu) + \max_{x \in K} |f(x)|)$. Therefore, the right-hand side of (24) is bounded above by³⁵

$$2(1 + \delta) \int_{\Omega(\Lambda)} \left(Z(\mu) + \max_{x \in K} f(x) - \min_{x \in K} f(x) \right) d\hat{\mu}(f).$$

We will show that

$$\sup_{\mu_t: t \geq 0} \int_{\Omega(\Lambda)} \left(Z(\mu) + \max_{x \in K} f(x) - \min_{x \in K} f(x) \right) d\mu_t(f)$$

converges to zero as Λ goes to infinity. For this, it suffices to show the convergence for

$$\sup_{\mu_t: t \geq 0} \int_{\Omega(\Lambda)} \left(\max_{x \in K} f(x) \right) d\mu_t(f),$$

since the other two terms can be treated similarly.³⁶ We will establish a stronger result, whose proof is in Appendix 9.2. Let $\mathcal{P}(K)$ denote the set of all finite partitions of K and, for each $\Pi \in \mathcal{P}(K)$ and $\bar{Z} \geq 0$, let $\mu_{\Pi}^{\bar{Z}}$ denote the probability measure over Θ corresponding to Brownian bridges with endpoints at consecutive elements of Π and endpoint values identically equal to \bar{Z} .

LEMMA 5 *For any constant $\bar{Z} \geq 0$,*

$$\lim_{\Lambda \rightarrow +\infty} \left\{ \sup_{\Pi \in \mathcal{P}(K)} \left\{ \int_{\Omega(\Lambda)} \left(\max_{x \in K} f(x) \right) d\mu_{\Pi}^{\bar{Z}}(f) \right\} \right\} = 0.$$

³⁵The inequality relies on the fact that, since $f(0) = 0$, $\max_{x \in K} f(x) \geq 0$ and $\min_{x \in K} f(x) \leq 0$.

³⁶The last term obtains by symmetry, the first term with the constant $Z(\mu)$ can in fact be incorporated in \bar{Z} in the argument following Lemma 5.

For each μ_t , $\max_{x \in K} f(x)$ is clearly dominated, in the sense of first-order stochastic dominance, by the same maximum under the distribution μ_Π , whose partition corresponds to the units of μ_t , and whose endpoints are equal to $Z(\mu)$, which is greater than $Z(\mu_t)$.³⁷ Applying Lemma 5 to $\bar{Z} = Z(\mu)$ thus proves the desired uniform convergence.

This implies that there exists, for any $\varepsilon > 0$, a positive threshold $\Lambda(\varepsilon)$ such that $|u^\Lambda(x, \mu_t, z) - u(x, \mu_t, z)| < \varepsilon$ for all $(x, t, z) \in K \times \mathbb{N} \times [0, Z(\mu)]$ and $\Lambda > \Lambda(\varepsilon)$. Therefore, (23) implies that, for a sequence converging to x , we have

$$\begin{aligned} u^\Lambda(x_{t_k}^*, \mu_{t_k}, Z(\mu_{t_k})) &\geq u(x_{t_k}^*, \mu_{t_k}, Z(\mu_{t_k})) - \varepsilon \\ &\geq u(x', \mu_{t_k}, Z(\mu_{t_k})) - \varepsilon \\ &\geq u^\Lambda(x', \mu_{t_k}, Z(\mu_{t_k})) - 2\varepsilon. \end{aligned}$$

Taking the limit as Λ goes to infinity, and using Lemma 4, we obtain that

$$u(x, \mu, Z(\mu)) \geq u(x', \mu, Z(\mu)) - 2\varepsilon.$$

Since ε was arbitrary, this proves that $u(x, \mu, Z(\mu)) \geq u(x', \mu, Z(\mu))$. Proposition 2 also implies that $u(x_{t_k}, \mu_{t_k}, Z(\mu_{t_k})) > (1 + \delta)Z(\mu_{t_k})$, which shows that $U(\mu, Z(\mu)) = u(x, \mu, Z(\mu)) \geq (1 + \delta)Z(\mu)$. Moreover, $f(x) \leq Z(\mu)$, since $Z(\mu) = \sup_t \{f(x_t)\}$, f is continuous, and x is a limit point of $\{x_t\}$, and $u(x, \mu, Z(\mu)) = f(x) + \delta Z(\mu) \leq (1 + \delta)Z(\mu)$, where the equality holds because $f(x)$ is known given μ . Therefore, $U(\mu, Z(\mu)) = (1 + \delta)Z(\mu)$. In particular, $f(x) = Z(\mu)$. ■

Proposition 8 and its proof also show that the value of marginal innovation converges to zero over time: for any x , $u(x, \mu_t, Z(\mu_t)) - (1 + \delta)Z(\mu_t)$ becomes nonpositive.

Step 4: Convergence to a Single Technology.

We prove three lemmas, from which uniqueness of the limit easily follows. For notational simplicity, we assume without loss of generality, for the remaining of the proof, that $K = [0, 1]$.

LEMMA 6 *Consider a unit u with extremities $x_l < x_r$ such that $0 < x_l$ and $x_r < 1$. Almost surely, either $[0, x_l]$ or $[x_r, 1]$ is visited finitely many times.*

³⁷Indeed, μ_Π is obtained from μ_t by raising to the level $Z(\mu)$ the payoff of each technology explored by time t .

Proof. Consider the innovation problem with initial domain $[0, x_l]$ (i.e., if only that technological domain were available), and let Z_l denote the limit value of the best explored technology, which we have shown to exist in Step 2. Similarly, we can define Z_r over $[x_r, 1]$. Z_l and Z_r are two random variables and $Z_l \neq Z_r$ a.s. Suppose, contradiction, that both $[0, x_l]$ and $[x_r, 1]$ were visited infinitely often. The sequence of technology choices on each of these intervals coincides with the sequence of equilibrium technological choices that would occur if the technological domain were limited to only that interval. Thus, the maxima reached on each subinterval in the more general problem must equal Z_l and Z_r . Since $Z_l \neq Z_r$ a.s., an agent will eventually prefer one subinterval over the other, which leads to a contradiction. ■

LEMMA 7 *For any history h and unit u with endpoints $x_l < x_r$ arising along that history, there exists a subinterval $I' = [x'_l, x'_r] \subseteq I = [x_l, x_r]$ with $x'_l < x'_r$ such that I' is visited finitely many times.*

Proof. Suppose, on the contrary, that for any partition of $I = [x_l, x_r]$, every element of the partition is visited infinitely often. Thus, every element of the partition must also contain an accumulation point from $\mathcal{M}(\{x_t\})$. Proposition 8 then implies that the best limiting payoff in every element of the partition is given by $Z(\mu)$. Since the choice of the partition was arbitrary, that can only happen for a set of paths of the Brownian bridge that has zero probability. ■

LEMMA 8 *For any $x \in \mathcal{M}(\{x_t\})$ and $\varepsilon > 0$, both the left and right neighborhood of x of length ε are visited infinitely often.*

Proof. Let

$$w = \liminf \{x_t : x_t > x, t \geq 0\}$$

and suppose, by contradiction, that $w > x$. For ε arbitrarily small, there exists T such that $x_T \in [w, w + \varepsilon)$. Let u denote the unit with interval $[x, x_T]$. Proposition 2, combined with the fact that $f(x) = Z(\mu)$, implies that $\gamma(u, Z(\mu)) > Z(\mu)$. Therefore, that there exists some $\tilde{x} \in (x, x_T)$ such that $u(\tilde{x}, \mu, Z(\mu)) > (1 + \delta)Z(\mu) = U(\mu, Z(\mu))$, a contradiction. ■

To conclude the proof, suppose by way of contradiction that there exist $x, y \in \mathcal{M}(\{x_t\})$, with $x < y$, and consider the interval $[x, y]$. By Lemma 8, we can always find two known actions x' and y' such that $x < x' < y' < y$. By Lemma 7, there exists a subinterval $[x'', y''] \subseteq [x', y']$ of positive measure that is visited only finitely many times. Lemma 6 then implies that

either $[0, x'']$ or $[y'', 1]$ is visited only finitely many times, which contradicts the hypothesis that both x and y are accumulation points.

9.1 Proof of Lemma 4

Let $\tilde{\Theta} = \Theta(\Lambda)$ (Λ is fixed throughout, so there is no ambiguity about the underlying space). Let $\{(x_n, \mu_n, z_n)\}$ be a sequence from $K \times \Delta(\tilde{\Theta}) \times [-\Lambda, \Lambda]$ which converges to $(x, \mu, z) \in K \times \Delta(\tilde{\Theta}) \times [-\Lambda, \Lambda]$. Then,

$$\begin{aligned} |v(x_n, \mu_n, z_n) - v(x, \mu, z)| &= \left| \int_{\tilde{\Theta}} r(g(x_n), z_n) d\mu_n - \int_{\tilde{\Theta}} r(g(x), z) d\mu \right| \\ &\leq \left| \int_{\tilde{\Theta}} [r(g(x_n), z_n) - r(g(x_n), z)] d\mu_n \right| + \left| \int_{\tilde{\Theta}} [r(g(x_n), z) - r(g(x), z)] d\mu \right| \\ &\quad + \left| \int_{\tilde{\Theta}} [r(g(x_n), z) - r(g(x), z)] d\mu_n \right| + \left| \int_{\tilde{\Theta}} r(g(x), z) d\mu_n - \int_{\tilde{\Theta}} r(g(x_n), z) d\mu \right| \\ &\leq \delta |z_n - z| + 2 \int_{\tilde{\Theta}} |r(g(x_n), z) - r(g(x), z)| d\mu \\ &\quad + \int_{\tilde{\Theta}} |r(g(x_n), z) - r(g(x), z)| d\mu_n + \left| \int_{\tilde{\Theta}} r(g(x), z) d\mu_n - \int_{\tilde{\Theta}} r(g(x), z) d\mu \right| \end{aligned}$$

The last term converges to zero by weak convergence of the beliefs. We focus on the second term

$$\int_{\tilde{\Theta}} |r(g(x_n), z) - r(g(x), z)| d\mu \leq (1 + \delta) \int_{\tilde{\Theta}} |g(x_n) - g(x)| d\mu$$

which converges to zero by the Bounded Convergence theorem. Next,

$$\int_{\tilde{\Theta}} |r(g(x_n), z) - r(g(x), z)| d\mu_n \leq (1 + \delta) \int_{\tilde{\Theta}} |g(x_n) - g(x)| d\mu_n$$

K is compact and every $g \in \tilde{\Theta}$ is continuous, hence also uniformly continuous. Fix $\varepsilon > 0$ and let

$$A\left(\frac{1}{m}, \varepsilon\right) = \left\{ g \in \tilde{\Theta} : \exists \lambda > \frac{1}{m} \text{ s.t. } |x - y| < \lambda \implies |g(x) - g(y)| < \varepsilon \right\}$$

By the previous observations, it also follows that for any $g \in \tilde{\Theta}$, there exists $m = m(g)$ such that $g \in A\left(\frac{1}{m'}, \varepsilon\right)$, $\forall m' > m$. Thus, $\tilde{\Theta} = \bigcup_{m=1}^{\infty} A\left(\frac{1}{m}, \varepsilon\right)$.

Next, since $\{A\left(\frac{1}{m}, \varepsilon\right)\}$ converges to $\tilde{\Theta}$, it follows that for any $\eta > 0$, there exists $M > 0$ such that $\mu\left(A\left(\frac{1}{m}, \varepsilon\right)\right) > 1 - \frac{\eta}{2}$, $\forall m > M$. Fix $\tilde{m} > M$, by weak convergence of beliefs, there exists $N > 0$ such that $|\mu_n\left(A\left(\frac{1}{\tilde{m}}, \varepsilon\right)^c\right) - \mu\left(A\left(\frac{1}{\tilde{m}}, \varepsilon\right)^c\right)| < \frac{\eta}{2}$, for any $n > N$.

Since $x_n \rightarrow x$, there exists $N' > N$ such that $|x_n - x| < \frac{1}{\tilde{m}}$, for any $n > N'$. Finally, we obtain, for $n > N'$,

$$\begin{aligned} \int_{\bar{\Theta}} |g(x_n) - g(x)| d\mu_n &= \int_{A(\frac{1}{\tilde{m}}, \varepsilon)} |g(x_n) - g(x)| d\mu_n + \int_{A(\frac{1}{\tilde{m}}, \varepsilon)^c} |g(x_n) - g(x)| d\mu_n \\ &\leq \sup_{g \in A(\frac{1}{\tilde{m}}, \varepsilon)} |g(x_n) - g(x)| + 2\Lambda \mu_n \left(A \left(\frac{1}{\tilde{m}}, \varepsilon \right)^c \right) \\ &\leq \varepsilon + 2\Lambda \left[\left| \mu_n \left(A \left(\frac{1}{\tilde{m}}, \varepsilon \right)^c \right) - \mu \left(A \left(\frac{1}{\tilde{m}}, \varepsilon \right)^c \right) \right| + \left| \mu \left(A \left(\frac{1}{\tilde{m}}, \varepsilon \right)^c \right) \right| \right] \\ &\leq \varepsilon + 2\Lambda \eta \end{aligned}$$

Since ε and η were arbitrary, this completes the proof. ■

9.2 Proof of Lemma 5

Letting $\Omega_+(\Lambda) = \{f : \max_{x \in K} f(x) > \Lambda - 1\}$ and $\Omega_-(\Lambda) = \{f : \min_{x \in K} f(x) < -(\Lambda - 1)\}$, we have $\Omega(\Lambda) = \Omega_+(\Lambda) \cup \Omega_-(\Lambda)$. Since $\bar{Z} \geq 0$, $\max_{x \in K} f(x)$ is nonnegative for all f in the support of any $\mu_\Pi \in \mathcal{P}(K)$. Therefore,

$$\int_{\Omega(\Lambda)} \max_{x \in K} f(x) d\mu_\Pi(f) \leq \int_{\Omega_+(\Lambda)} \max_{x \in K} f(x) d\mu_\Pi(f) + \int_{\Omega_-(\Lambda)} \max_{x \in K} f(x) d\mu_\Pi(f).$$

We will prove that the first term (the harder one) converges to zero as $\Lambda \rightarrow \infty$, uniformly in Π . The second term can be treated similarly.

Put in the language of probability theory, we need to show that, for each $\bar{Z} \geq 0$, the family of random variables $\{X_\Pi = \max_{x \in K} f(x) : f \sim \mu_\Pi\}_{\Pi \in \mathcal{P}(K)}$ is uniformly integrable.³⁸ To show uniform integrability, it suffices to prove that there exists $p > 1$ such that.³⁹

$$\sup_{\Pi \in \mathcal{P}(K)} E[X_\Pi^p] < \infty$$

Without loss of generality, we set $\bar{Z} = 0$ (other cases follow by translation) and $K = [0, 1]$ (other cases follow by the scaling property of Brownian motion). For each Π , we have

$$Pr(X_\Pi \leq \Lambda) = \prod_{\pi_i \in \Pi} Pr(X_i \leq \Lambda),$$

³⁸A family $\{X_i\}_{i \in I}$ is uniformly integrable if $\lim_{\Lambda \rightarrow +\infty} \{\sup_i \{E[|X_i| : |X_i| > \Lambda]\}\} = 0$.

³⁹See, e.g., Durrett (1996), Exercise 4.5.1., p. 260.

where $\{\pi_i\}_i$ describes the units of the partition Π , X_i is the maximum of f over π_i , and we are using the fact that the variables $\{X_i\}_i$ are independently distributed. Moreover, each X_i is the maximum of a Brownian bridge with width δ_i (the width of π) and endpoints equal to 0. This implies that⁴⁰

$$Pr(X_i \leq \Lambda) = 1 - e^{-2\Lambda^2/\delta_i^2}.$$

Therefore, we can compute the density of X_Π , and obtain, for $p = 2$,

$$E[X_\Pi^2] = \int_{\mathbb{R}_+} \Lambda^2 \sum_i \left(\prod_{j \neq i} \left(1 - e^{-2\Lambda^2/\delta_j^2} \right) \right) e^{-2\Lambda^2/\delta_i^2} \frac{4\Lambda}{\delta_i^2} d\Lambda.$$

For each i , the product with respect to j is bounded by 1, implying that

$$E[X_\Pi^2] \leq 4 \sum_i \int_{\mathbb{R}_+} \frac{\Lambda^3 e^{-2\Lambda^2/\delta_i^2}}{\delta_i^2} d\Lambda.$$

Making, for each i , the change of variable $u_i = \Lambda/\delta_i$, we obtain

$$E[X_\Pi^2] \leq 4 \sum_i \int_{\mathbb{R}_+} \delta_i^2 u_i^3 e^{-2u_i^2} du_i.$$

Since, for any partition Π , $\sum_i \delta_i^2 \leq \sum_i \delta_i = 1$, we conclude that

$$\sup_{\Pi \in \mathcal{P}(K)} E[X_\Pi^2] \leq C$$

where $C = 4 \int_{\mathbb{R}_+} u^3 e^{-2u^2} du < \infty$.

10 Proof of Theorem 3

We start by showing the existence of a threshold $\tilde{\rho}$ above which $y^R(\rho) = 0$. Suppose, first, that $\kappa \leq c'(0)$, so that an old agent chooses the best explored technology (as observed in the

⁴⁰See, e.g., Durrett (1996), Exercise 7.8.1., p. 433. The formula given there is for a Brownian bridge with width equal to 1. The general formula obtains by the scaling property of Brownian motion, which is easily shown to be inherited by the Brownian bridge: letting $\{B_t^a\}_{t \in [0, a]}$ denote a Brownian bridge with endpoints 0 and a and endpoint values equal to 0, B_t^a has the law of $B_t - \frac{t}{a} B_a$, where B is the standard Brownian motion.

main text). The expected utility of a young agent from choosing technology $x > \bar{x}$ is

$$f(\bar{x}) + \kappa y - c(y) + \delta \left\{ (f(\bar{x}) + \kappa y) \left(1 - \Phi \left(\frac{\rho - \kappa y}{\sigma \sqrt{y}} \right) \right) + \sigma \sqrt{y} \phi \left(\frac{\rho - \kappa y}{\sigma \sqrt{y}} \right) \right. \quad (25)$$

$$\left. + (\rho + f(\bar{x})) \Phi \left(\frac{\rho - \kappa y}{\sigma \sqrt{y}} \right) \right\} \quad (26)$$

$$= (1 + \delta)[f(\bar{x}) + \kappa y] - c(y) + \delta \left\{ (\rho - \kappa y) \Phi \left(\frac{\rho - \kappa y}{\sigma \sqrt{y}} \right) + \sigma \sqrt{y} \phi \left(\frac{\rho - \kappa y}{\sigma \sqrt{y}} \right) \right\}, \quad (27)$$

where $\rho = z - f(\bar{x})$. The first-order condition is

$$\kappa \left(1 + \delta - \delta \Phi \left(\frac{\rho - \kappa y}{\sigma \sqrt{y}} \right) \right) + \frac{\delta \sigma}{2\sqrt{y}} \phi \left(\frac{\rho - \kappa y}{\sigma \sqrt{y}} \right) = c'(y). \quad (28)$$

The left-hand side of (28) approaches κ as ρ increases and, for fixed ρ , the left-hand side converges to $\kappa(1 + \delta)$ as $y \rightarrow +\infty$, and to κ as $y \rightarrow 0$. Also, the left-hand side is bounded above by

$$\kappa(1 + \delta) + \frac{\delta \sigma^2 \phi(1)}{2\rho}$$

which converges to $\kappa(1 + \delta)$, as ρ increases. Thus, $\lim_{y \rightarrow +\infty} c'(y) > \kappa(1 + \delta)$ (and $c''(0) > 0$, if $\kappa = c'(0)$) implies that there exists $\tilde{\rho} > 0$ such that $\rho > \tilde{\rho}$ implies $y^R(\rho) = 0$.

Differentiating (28) with respect to ρ yields

$$-\frac{\delta \phi}{2\sigma \sqrt{y}} \left[\kappa + \frac{\rho}{y} \right] < 0$$

for all $\rho, y > 0$. This implies that $y^R(0) \geq y^R(\rho)$ for all $\rho > 0$. Repeating the argument used to prove Theorem 2, we conclude that radical innovation ends in finite time a.s.

Suppose now that $\kappa > c'(0)$. An old agent experiments radically with a size equal to $y^{O,R} > 0$ if and only if $\rho \leq \xi$, as shown in the main text. We assume without loss of generality that $\rho_t > \xi$, so that an old agent does not innovate today and then the expected utility today of any $x > \bar{x}_t$ for a young agent is simply

$$\begin{aligned} & E[f(x) - c(x - \bar{x}_t)] + \delta \{ E[f(x) + \xi | f(x) \geq z_t - \xi] \text{Prob}(f(x) \geq z_t - \xi) + z_t \text{Prob}(f(x) < z_t - \xi) \} \\ &= (1 + \delta)(f(\bar{x}_t) + \kappa y) - c(y) + \delta \left\{ \sigma \sqrt{y} \phi \left(\frac{\rho_t - \xi - \kappa y}{\sigma \sqrt{y}} \right) + \xi + (\rho_t - \xi - \kappa y) \Phi \left(\frac{\rho_t - \xi - \kappa y}{\sigma \sqrt{y}} \right) \right\} \end{aligned}$$

The first-order condition is

$$\kappa \left[1 + \delta - \delta \Phi \left(\frac{\rho_t - \xi - \kappa y}{\sigma \sqrt{y}} \right) \right] + \frac{\delta \sigma}{2\sqrt{y}} \phi \left(\frac{\rho_t - \xi - \kappa y}{\sigma \sqrt{y}} \right) = c'(y) \quad (29)$$

The right-hand side is always greater than or equal to κ . Since $\rho_t - \xi > 0$, the right-hand side converges to κ as $y \rightarrow 0$, and to $\kappa(1 + \delta)$ as $y \rightarrow +\infty$. Since also $c'(0) < \kappa$, there must exist an interior solution to the first-order equation (29). When $\lim_{y \rightarrow +\infty} c'(y) > \kappa(1 + \delta)$, the solution is unique for high values of ρ because the left-hand side of (29) approaches κ pointwise as ρ becomes arbitrarily large. As ρ increases, $y^R(\rho)$ approaches y^{OR} : the optimal size of innovation for a young agent converges to the optimal size for an old agent. This follows from the fact that the first-order condition becomes, dropping negligible terms, $\kappa \approx c'(y)$. The maximized expected utility is approximately equal to

$$\begin{aligned} & (1 + \delta)(f(\bar{x}_t) + \xi) + \delta \left[\kappa y^{O,R} \left(1 - \Phi \left(\frac{\rho_t - \xi - \kappa y^{O,R}}{\sigma \sqrt{y^{O,R}}} \right) \right) \right. \\ & \left. + \sigma \sqrt{y^{O,R}} \phi \left(\frac{\rho_t - \xi - \kappa y^{O,R}}{\sigma \sqrt{y^{O,R}}} \right) + (\rho_t - \xi) \Phi \left(\frac{\rho_t - \xi - \kappa y^{O,R}}{\sigma \sqrt{y^{O,R}}} \right) \right] \\ & \approx f(\bar{x}_t) + \xi + \delta z_t \end{aligned}$$

Since we assumed that $\rho_t > \xi$, it follows that $f(\bar{x}_t) + \xi + \delta z_t < (1 + \delta)z_t$. The only candidate for radical innovation gives an expected outcome which is lower than what the young agent could get by simply exploiting. Thus, there exists $\tilde{\rho} > 0$ such that the young agent prefers exploitation for any gap greater than $\tilde{\rho}$.

In order to replicate the steps used to prove Theorem 2, we still need to show that there is an upper bound on the equilibrium size of innovation. When $\rho > \xi$, the right-hand side of (29) is strictly decreasing in ρ . Thus, the unique positive solution of the first-order condition when $\rho = \xi$ provides the desired upper bound. When $\rho \leq \xi$, the old agent is experimenting with a fixed size $y^{O,R}$ (independent of the gap). Replicating the argument for $\kappa = 0$, one may show that, on the range $\rho \in [0, \xi]$, the value and size of radical innovation are decreasing in ρ , providing an upper bound on the size of radical innovation for a young agent, uniform over all histories.

References

- Acemoglu, D., 2008, *Introduction to Modern Economic Growth*. Princeton University Press, Princeton, NJ.
- Acemoglu, D., M. Dahleh, I. Lobel, and A. Ozdaglar, 2011, “Bayesian Learning in Social Networks,” *Review of Economic Studies*, 78, 1201–1236.
- Acemoglu, D., G. Gancia, and F. Zilibotti, 2010, “Competing Engines of Growth: Innovation and Standardization,” *Journal of Economic Theory*, forthcoming.
- Aghion, P., P. Bolton, C. Harris, and B. Jullien, 1991, “Optimal Learning by Experimentation,” *Review of Economic Studies*, 58, 621–654.
- Bala, V., and S. Goyal, 1998, “Learning from Neighbours,” *Review of Economic Studies*, 65, 595–621.
- Banerjee, A., 1992, “A Simple Model of Herd Behavior,” *Quarterly Journal of Economics*, 107, 797–817.
- Bikhchandani, S., D. Hirshleifer, and I. Welch, 1992, “A Theory of Fads, Fashion, Custom, and Cultural Change as Information Cascades,” *Journal of Political Economy*, 100, 992–1026.
- Billingsley, P., 1968, *Convergence of Probability Measures*. John Wiley and Sons, New York, NY.
- Bramoullé, Y., and G. Saint-Paul, 2010, “Research Cycles,” *Journal of Economic Theory*, 145, 1890–1920.
- Callander, S., 2011, “Searching and Learning by Trial-and-Error,” *American Economic Review*, 101, 2277–2308.
- Dembo, A., and O. Zeitouni, 1998, *Large Deviations Techniques and Applications*. Springer-Verlag, New York, NY.
- Durrett, R., 1996, *Probability: Theory and Examples*. Duxbury Press, 2nd ed.
- Easley, D., and N. Kiefer, 1988, “Controlling a Stochastic Process with Unknown Parameters,” *Econometrica*, 56, 1045–1064.

- Edlin, A., and C. Shannon, 1998, “Strict Monotonicity in Comparative Statics,” *Journal of Economic Theory*, 81, 201–219.
- Garicano, L., and E. Rossi-Hansberg, 2009, “Organizing Growth,” *Journal of Economic Theory*, forthcoming.
- Griliches, Z., 1992, “The Search for R&D Spillovers,” *Scandinavian Journal of Economics*, 94, S29–S47.
- Hall, B., 1996, “The Private and Social Returns to Research and Development,” *Technology, R&D, and the Economy*, Brookings Institution and the American Enterprise Institute, Washington, DC, 140–83.
- Jones, C., and J. Williams, 1998, “Measuring the Social Return to R&D,” *Quarterly Journal of Economics*, 113, 1119–1135.
- Jovanovic, B., and Y. Nyarko, 1996, “Learning by Doing and the Choice of Technology,” *Econometrica*, 64, 1299–1310.
- Jovanovic, B., and R. Rob, 1990, “Long Waves and Short Waves: Growth Through Intensive and Extensive Search,” *Econometrica*, 58, 1391–1409.
- Karatzas, I., and S. Shreve, 1991, *Brownian Motion and Stochastic Calculus*. Springer-Verlag, New York, NY.
- Kauffman, S., L. Lobo, and W. Macready, 2000, “Optimal Search on a Technology Landscape,” *Journal of Economic Behavior & Organization*, 43, 141–166.
- Lee, I., 1993, “On the Convergence of Informational Cascades,” *Journal of Economic Theory*, 61, 395–411.
- Mandler, M., 2011, “Maximizing the Value of Science: No News can be Good News,” *mimeo*.
- Matsuyama, K., 1999, “Growing Through Cycles,” *Econometrica*, 67, 335–347.
- McLennan, A., 1984, “Price Dispersion and Incomplete Learning in the Long Run,” *Journal of Economic Dynamics and Control*, 7, 331–347.
- Milgrom, P., and I. Segal, 2002, “Envelope Theorems for Arbitrary Choice Sets,” *Econometrica*, 2, 583–601.

Milgrom, P., and C. Shannon, 1994, “Monotone Comparative Statics,” *Econometrica*, 62, 157–180.

Niehaus, P., 2011, “Filtered Social Learning,” *Journal of Political Economy*, 119, 686–720.

Rothschild, M., 1974, “A Two-Armed Bandit Theory of Market Pricing,” *Journal of Economic Theory*, 9, 185–202.