Social Experimentation and Innovation Cycles

Umberto Garfagnini
ITAM School of Business

Bruno Strulovici
Northwestern University

July 15, 2013

Abstract

This paper studies social learning by a succession of forward-looking agents who endogenously choose their information according to the exploration-exploitation trade-off at the heart of the experimentation paradigm. We analyze social experimentation with a continuum of interdependent technologies, emphasizing the distinction, fundamental in the literature on innovations, between radical and marginal innovations. We characterize alternating cycles between radical and marginal innovations, and find that radical innovation must stagnate in the long run for all parameters of our model. We also establish a negative relationship between past innovation successes and the *magnitude* of radical innovations.

Keywords: Social Learning, Experimentation, Innovation, R&D, Stagnation.

JEL Codes: C73, D83, O3

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. Strulovici gratefully acknowledges financial support from the NSF under Grant No. 1151410 and the Alfred P. Sloan Foundation. Email addresses: umberto.garfagnini@itam.mx and b-strulovici@northwestern.edu.

1 Introduction

Modeling the dynamics of technological innovation is challenging: one must include research and experimentation by individual firms, the transmission of knowledge across firms, and the fact that the set of available technologies is time dependent, virtually unlimited, and entails highly uncertain payoffs. All these aspects are necessary, for instance, to study whether radical innovation is self-sustainable in the long run or requires external intervention, a question that has received longstanding interest from economists and policymakers alike.[1]

One source of inspiration to approach this question comes from social learning models which have served since Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992), as the main paradigm to analyze knowledge accumulation by successive agents, from the adoption of new technologies to cultural change. Those models examine the asymptotic efficiency and fragility of social learning under restrictive assumptions (myopic agents, exogenous signals, stationary and limited choice sets), however, which make them ill-suited to study important knowledge accumulation processes, such as technological innovation.

This paper draws on the experimentation, social learning, and innovation literatures to analyze the dynamics of radical and marginal innovations. Unlike social learning (or "herding") models, where each agent learns only from the actions of his predecessors and an exogenous signal, we model firms as two-period agents, with an incentive to explore new technologies in their "young" period and exploit existing ones in their "old" period. Compared to strategic experimentation models, which typically consider tradeoffs between a pre-specified, single risky technology and a safe one whose payoff is perfectly known,[2] we model technologies as a continuum of related "arms" of varying risks and costs. This spatial representation allows us to make the distinction, crucial in the industrial organization literature, between radical and marginal innovation. We define marginal innovations as convex combinations of existing technologies, which bear no cost other than the opportunity cost of forgoing the exploitation of known technologies. By contrast, radical innovations are technologies that lie beyond the convex hull of existing technologies, and incur an additional investment cost.

[1]See, e.g., Griliches (1992), Hall (1996), Gordon (2012). Recent press coverage includes Slywotsky (2009), Zakaria (2011). Policymakers' interest was highlighted by President Obama's State of the Union Address in 2011.

[2]Main references include Bolton and Harris (1999), Keller, Rady, and Cripps (2005), and Keller and Rady (2010).

Of central importance, in our analysis, is the interplay between the outcomes of past innovations and incentives for future innovation. The question is not only whether radical innovations take place, but also *how radical*, or ambitious, those innovations are. We find that past successes reduce not only the value of radical innovation, but also its incremental value.[3] As a result, past successes result in less frequent and less ambitious innovations.

To compare the values of radical and marginal innovations, we partition the technological space according to *units*, i.e., intervals whose interior consists of unexplored technologies and endpoints consist of explored ones. The partition gets finer over time, as more technologies get explored. The value of innovation of a given unit depends on its width and on the payoffs of the technologies that delimit it. In addition, it may depend on technologies outside of the unit *only* through the highest known payoff. This quasi-independence from outside technologies results from the assumption that payoffs are drawn from a Brownian motion with drift: the conditional payoff distribution within any unit, given the payoffs at its endpoints, is independent from all technologies outside of the interval. Thus, while our model allows for payoff correlation across technologies, the set of feasible technologies can be partitioned, after any finite history, into units that are similar to the independent arms arising in the standard multi-armed bandit problem, up to the aforementioned dependence on the highest known payoff.

Standard intuition from the social learning literature would suggest that whether radical innovation is sustained in the long run might depend on the parameters of the model and/or on the particular realization of the signals (here, payoffs) received by the agents. Could the drift of the Brownian motion, which determines the expected payoff of new technologies, influence the answer? Surprisingly, however, we find that radical innovation stagnates for all histories, regardless of the drift and cost of radical innovation. There is no fragility in this result: whether big or small, shocks do not affect long-run stagnation. The intuition for the result may be summarized as follows: if recent radical innovations have been disappointing relative to known technologies, which must eventually happen with probability 1, then because such radical innovations serve as the basis of further radical innovation, agents prefer to stay within the confines of marginal innovation. Therefore, a form of informational cascade arises, in which no agent accepts to bear the cost of radical innovation, and the outlook for further radical innovation is frozen and negative.

[3] To avoid confusion with "marginal" innovation, we use the term "incremental" to describe derivatives of benefits and cost functions.

The dynamics of innovation may be more fully described as follows. While marginal innovation takes place, the informational ("option") value of marginal innovation decreases, and this may spur a new cycle of radical innovations. As new technologies become available, however, the option value of marginal innovation gets replenished, making marginal innovation more attractive, and radical innovation gives again way to marginal innovation, generating a new innovation cycle. Eventually, known technologies prevail over radical innovation, which is then abandoned forever.

Our paper thus characterizes the short and long run behavior of innovation, showing the emergence of finitely many cycles of radical and marginal innovations. We also show that innovations and beliefs have a well-defined limit, despite the fact that both the technology and payoff spaces are non compact. Showing this innocuous-looking result is actually challenging. One technical contribution of our paper is to extend to an unbounded domain (both for the *domain* of technologies and the *range* of their payoffs) the techniques developed by Easley and Kiefer (1988) to show the existence of a well-defined limit for agents' beliefs resulting from experimentation.

The papers closest to ours are Jovanovic and Rob (1990) and Callander (2011). Our model shares with those works the use of Brownian motion as the source of payoff uncertainty across related technologies (or "policies"). In those models, however, agents live only for one period. Therefore, the trade-off between exploration and exploitation, at the heart of the experimentation literature and of our paper, is absent. Jovanovic and Rob circumvent this issue by allowing the agent to learn, at some fixed cost, the value of a new "technique," and then decide whether to try the new technique or use an old one. Thus, by disentangling information acquisition from technological choice, that model restores some learning from otherwise myopic agents. As a result of this and other modeling assumptions, if an agent chooses a new technique, at any time, then all previous techniques are forever abandoned. Another consequence is that the size of radical innovation (moving to a new technique) is independent of past technologies. Our model provides arguably richer and more sophisticated innovation dynamics than Jovanovic and Rob's seminal model of innovation over Brownian paths, and is closer to the spirit of experimentation.[4] Callander (2011) takes a different

[4] In a short extension, Jovanovic and Rob consider the case of two-period agents. That extension contains one result: for a given common history, if a one-period agent is indifferent between old and new technologies, a two-period agent will strictly prefer the new one. That result is hard to interpret, however, because agents with different life spans typically generate different histories. The short and long run dynamics of innovation

approach: instead of looking for the maximum payoff technology or "policy," his agents want to find an ideal policy, whose value is normalized to zero. The goal of those agents is therefore to find some zero of a Brownian path. Even a myopic agent is willing to try a new technology, because interpolation between negative and positive values yields an expectation that is closer to his ideal policy.[5]

The paper is organized as follows. Section 2 describes the baseline model. Section 3 introduces the concepts of values of radical and marginal innovation. Section 4 investigates the short-run dynamics of innovation, represented by innovation cycles, and the main stagnation result. Section 5 compares equilibrium innovation dynamics with the socially efficient path. Section 6 shows that the stagnation result persists in the case of optimistic beliefs. Section 7 compares our mechanism of stagnation with those that have been presented in the literature on growth and innovation. All the proofs are in the Appendix, which also contains an extension to the case in which the cost of radical innovation depends on the outcome of past innovation, and a micro-foundation of the model.

2 The Model

We introduce an overlapping generations model with the following characteristics. An agent is born at each period $t \in \mathbb{N} = \{0, 1, \ldots, \}$, who lives for two periods, "young" and "old." The agent is risk neutral and chooses at each period a technology x in the technological space $X = [0, \infty)$. The payoff $f(x)$ of technology x is initially unknown, except at the origin where $f(0) = 0$.[6]

A young agent inherits from the contemporary old agent the knowledge of all technologies and payoffs that have previously been tried. We assume that this information transmission is costless and non-strategic.[7] The history h_t at time t consists of all technology-payoff pairs

are also left open for that case.

[5]Because agents are averse to risk, experimentation may still stall, as the benefit from interpolation may be dominated by the risk resulting from the positive variance of new policies. The dynamics of that paper are very interesting and also very different from ours.

[6]The function $f(\cdot)$ could represent utils instead of payoffs, as long as agents are expected-utility maximizers. One could also consider the larger domain $X = \mathbb{R}$. Radical innovation would then have "left" and "right" components.

[7]The older generation is not affected by the choice of the younger one and has therefore no incentive to manipulate the transmission. It could benefit from selling the information. We discuss patents in Section 7.

that have been experienced in the past. We let z_t denote the highest payoff among explored technologies, $\bar{x}_t$ denote the rightmost explored technology, i.e., the current *frontier* of the explored domain, and $g_t = z_t - f(\bar{x}_t) \geq 0$ denote the difference, or *gap*, between the payoffs of the best explored technology and the frontier technology. The set $[0, \bar{x}_t]$ of technologies that are convex combinations of previously explored ones is called the *active domain*.

The payoff function f is assumed to follow the distribution of a Brownian motion with drift κ and volatility $\sigma > 0$.[8] For simplicity we first analyze the case in which $\kappa = 0$, and then extend our main results to general drifts in Section 6. In the benchmark case, the payoff of a technology $x > \bar{x}_t$ has a normal distribution with mean $f(\bar{x}_t)$ and variance $(x - \bar{x}_t)\sigma^2$. Technologies to the right of $\bar{x}_t$ thus have the same expected payoff, and a variance that increases with their distance from $\bar{x}_t$.

The technology space may be partitioned according to previously tried technologies: the partition consists of the finitely many bounded intervals whose endpoints have been explored and whose interior points have not, and of the unbounded interval $[\bar{x}_t, +\infty)$. A *unit* is defined by any such interval along with the values of its endpoint payoffs. The distribution of payoffs within a given unit is described by a Brownian bridge: it is the distribution of Brownian motion on some interval with known end values.[9] In particular, it is conditionally independent of the payoff of all observed technology-payoff pairs outside of the unit.

For any bounded unit u, let m (M) denote the smaller (larger) of the two endpoint payoffs, L the width of the underlying interval, and $d = M - m \geq 0$ the difference between endpoint payoffs. These variables clearly depend on the unit they are attached to, but the reference to u is omitted when there is no ambiguity. The unbounded unit at time t is denoted u^∞, again omitting the reference to t when there is no ambiguity.

A technology x lying in some bounded unit u with endpoints $x_l < x_r$ has a normally distributed payoff with mean

$$f(x_l) + \frac{f(x_r) - f(x_l)}{x_r - x_l}(x - x_l) \tag{1}$$

and variance

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

[8] More precisely, we start with a filtered probability space $(\Omega, \mathcal{F}, \mathbb{P})$ satisfying the usual regularity conditions (see, e.g., Karatzas and Shreve 1991) and whose outcomes are identified with the paths of a Brownian motion.

[9] We refer the reader to Billingsley (1968, p. 64) for an introduction to Brownian bridges.

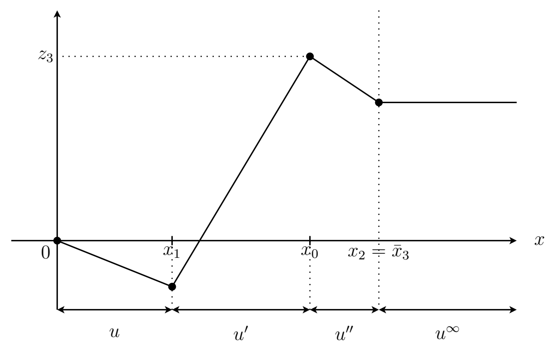

Figure 1: History of units after three periods.

The expected payoff of a technology increases linearly from the endpoint technology with the worst payoff to the one with the highest payoff. 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. Figure 1 represents an innovation history up to $t = 3$.

The literature on growth and technological change has taken numerous approaches to define marginal and radical innovations, which include informational and payoff components. Exploiting the spatial structure of our model, we define those concepts as follows.

A **marginal innovation** in period t is a technological choice $x \in [0, \bar{x}_t]$, and is assumed to be costless.[10] A **radical innovation** is a technology that lies beyond the frontier $\bar{x}_t$, and incurs a 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. Encompassing both cases, the cost of innovation is equal to $c(x - \bar{x})$, where c is twice continuously differentiable, increasing, convex, and such that i) $c(y) \leq 0$ for $y \leq 0$ (so that marginal innovation is costless) and ii) either $c'(0) > 0$ or $c''(0) > 0$.[11] A radical innovation therefore creates a positive *payoff* externality on future generations, in addition to an informational externality, because all technologies between the old and new frontiers become available at

[10]The 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.

[11]In an extension (Appendix E), we allow c to also depend on z, to capture the idea that current technologies affect the cost of further innovations.

no cost.

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 explored technology (recall that, in this benchmark model, radical innovations all have the same expected payoff as the frontier's payoff). A young agent solves the optimization problem

$$U(h_t) = \sup_{x \in X} E_{h_t} \left[f(x) - c(x - \bar{x}_t) + \delta \max\{f(x), z_t\} \right]. \tag{3}$$

3 Marginal and Radical Values of Innovation

The first step to characterize the dynamics of innovation is to address the correlation across payoff technologies: trying a technology reveals information not only about that technology, but also about nearby technologies, with an accuracy that decreases as one moves away from that technology.[12] Our strategy is to group technologies according to the units defined in the previous section and exploit the fact that experimenting with a given technology changes only the payoff distribution of technologies that belong to the same unit. The optimization problem of a young agent may be decomposed as, first, choosing one of finitely many units and, second, which technology to pick within that unit. The advantage of this decomposition is that we can characterize the value of each unit according to a simple index, as established by Theorem 1. This index determines the *value of innovation* of the unit which forms a key block in the analysis of innovation cycles, performed in the next section.

We fix a history $h = h_t$ with maximum explored payoff $z = z_t$ and gap $g = z - f(\bar{x})$.

THEOREM 1 *One can associate to each unit u a* **value of innovation**, $V(u, z)$, *which has the following properties:*

i) It is optimal for the young agent to choose a technology in the unit with the highest value of innovation.

[12]The relation is in fact more subtle: it also depends on how far other explored technologies are from the newly tried technology.

ii) There exist functions $\eta : \mathbb{R}_+^2 \to [1, \infty)$, and $\eta^\infty : \mathbb{R}_+ \to \mathbb{R}_+$ such that

$$V(u,z) = \begin{cases} (m-z) + d\,\eta\left(\frac{\sqrt{L}}{d}, \frac{z-m}{d}\right) & \text{if } u \neq u^\infty \\ \\ \eta^\infty(g) - g & \text{if } u = u^\infty \end{cases} \tag{4}$$

iii) If u is bounded and $M < z$, then $V(u,z)$ is increasing in L, m and M (other things equal).

Moreover, it is strictly optimal, for a bounded unit, to choose a technology closer to the endpoint with the higher payoff.[13]

The value of innovation of a unit represents the (normalized) highest lifetime expected payoff that a young agent can get, given history h_t, when his choice in the first period is restricted to a technology within the unit u *in excess of the payoff guaranteed by exploitation, z*. When his choice in the first period is restricted to technologies in u, a young agent's value function may be written as $(1+\delta)\,[z + V(u,z)]$.

Some units may have a negative value of innovation. Choosing from those units is always suboptimal: receiving the best explored payoff z in both periods dominates such choices. The next proposition shows a stronger result: the value of innovation of a unit is nonincreasing in z and, hence, over time. Once a unit gets a negative value of innovation, it is therefore abandoned forever.

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.*

This monotonicity property is shown in two steps. Consider, first, a unit u that does *not* contain the best explored technology, and suppose that the payoff of that technology is increased from z to $z' > z$. The increase has no effect on the payoff distribution of technologies inside u, but it reduces the probability that any technology in u beats the best explored technology. This, intuitively, reduces the value of innovating in that unit, implying that $V(u,z') < V(u,z)$. Now, consider a unit u whose endpoints include the best explored technology. For such a unit, a higher value of z increases, linearly, the expected payoff of all technologies inside the unit. The variance of payoffs within that unit is unaffected (this

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

is a standard property of Brownian bridge), however, so exploitation with the best explored technology is relatively more appealing than before the increase, which again reduces the value of innovation.[14]

A common issue in experimentation and social learning models is whether, and with what probability, agents converge in finite time to some specific action. The next result shows that *exploitation* – choosing a known technology – is strictly suboptimal, which implies that a new unit is created at each period and, hence, that the partitioning in units of the technology space becomes strictly finer over time.

PROPOSITION 2 (EXPLOITATION) *If u contains the best explored technology, then $V(u, z) > 0$.*

This proposition implies that choosing z for both periods is strictly dominated by choosing another technology included in a unit u whose endpoint payoffs include z. Intuitively, a slight departure from the best explored technology reduces the expected payoff of the agent, but it also creates an option value that dominates that reduction, because volatility increments of Brownian motion (of order $\sqrt{dt}$) dominate expectation increments (of order dt). By the same argument, exploitation would remain suboptimal *even if marginal innovation were costly*, as long as the cost of marginal innovation goes smoothly to zero as one gets closer to known technologies.

To distinguish waves of marginal and radical innovations, we now introduce the key concepts of the paper. The **value of marginal innovation**, $V^M(h_t)$, is defined as the maximum value of innovation over all finite units. The **value of radical innovation** is defined as the value of innovation of the unbounded unit: $V^R(h_t) = V(u^\infty(h_t), z_t)$. From Theorem 1, an agent prefers radical over marginal innovation if and only if the value of radical innovation exceeds the value of marginal innovation.

Theorem 1 further implies that the value of radical innovation depends on h_t only through the gap $g_t = z_t - f(\bar{x}_t)$: distinct histories that induce the same gap yield identical incentives

[14] A 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: $(f(x)-f(x_l))/(z-f(x_l)) \to_{z\to\infty} m + (z-m)(x-x_l)/(x_h-x_l)$ 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.

for radical innovation; in particular, the payoff z of the best explored technology does not matter per se. To express this particular dependence on history, we will denote by $V^R(g_t)$ the value of radical innovation at time t. The next result provides an analogue of Proposition 1 for the unbounded unit.

Proposition 3 (Value of Radical Innovation) *$V^R(g)$ is decreasing in g. Moreover, if $V^R(g_t)$ is negative at any time, radical innovation is abandoned forever after.*

Intuitively, incentives for radical innovation are driven by the probability of finding a new technology whose payoff exceeds z_t. Since the payoff distribution of new technologies is pinned down by the payoff at the frontier $\bar{x}_t$, a higher gap reduces the probability of this event. If the value of radical innovation becomes negative, it can never be positive again, because $\bar{x}_t$ remains frozen and z_t can only increase. This proposition thus has far-reaching consequences for the sustainability of radical innovation in the long run, which are exploited in Theorem 2.

4 Innovation Cycles

In our model, social experimentation takes the form of innovation cycles, characterized in this section, in which radical innovation is followed by marginal innovation. Those cycles capture the Schumpeterian pattern of successful innovations followed by imitation. Radical innovations generate a positive externality on all future generations, by expanding the active domain and, hence, the set of marginal innovations. While marginal innovation refines knowledge about technologies in the active domain, it is characterized by more predictable outcomes and 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. Following a highly successful radical innovation, the best explored technology lies at the frontier. In that case, further radical innovation is equally likely to outperform or underperform the payoff of the current frontier, and a wave of radical innovations can occur. This and other results are formalized in Proposition 4.

Proposition 4 (Innovation Cycles) *Suppose that, following some history h_t, $V^R(g_t) > 0$. Then, innovation has the following properties.*

- *The probability that radical innovation occurs at time* t *or at some future date is strictly positive, regardless of the value of marginal innovation.*
- *If radical innovation takes place at time* t*, there exist (history-dependent) cutoffs* f_t^R *and* f_t^M *such that, letting* f_t *and* u_t *respectively denote the payoff resulting from that radical innovation and the newly created unit (i.e., the interval between the old and new boundaries),*
  - *i) radical innovation is optimal at time* $t+1$ *if* $f_t > f_t^R$*, and suboptimal if* $g_t = 0$ *and* $f_t < f_t^R$*, and*
  - *ii) if marginal innovation is optimal at time* $t+1$*, it takes place in* u_t *if and only if* $f_t > f_t^M$ *and* $f_t^M < f_t^R$*.*

Thus, highly successful radical innovation begets more radical innovation, while moderately successful innovation is followed by marginal innovation. To appreciate the subtlety of the first result, we observe that a high payoff 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 payoff distribution on the newly created unit is, for very high payoff realizations, roughly a straight line, with very low variance.[15] This means that technologies in the new unit have a lower expectation than the last technology, and a low variance. By 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. Whether marginal innovation occurs within the newly created unit, however, depends on whether or not $f_t^M < f_t^R$. Otherwise, marginal innovation, when optimal, will take place in the old domain.

Our results thus far may be summarized as follows: from Proposition 2, we know that exploitation is never optimal. Thus, any wave of marginal innovations reduces the width of units in the active domain. This, all else equal, reduces the value of innovation in the active domain, by Theorem 1. The value of marginal innovation may then decrease to the point of triggering a new round of radical innovation, which continues until it leads to a disappointing payoff, triggering a new wave of marginal innovations to explore further the

[15] This result is well-known in the literature on large deviations, see e.g., Dembo and Zeitouni (1998). See also Footnote 14.

units created within the new frontier, giving rise to alternating cycles between radical and marginal innovations.

Because the technological domain is unbounded, it is a priori possible that agents indefinitely expand the boundary of the active domain as time goes by. Such qualitative property of the model would correspond to unlimited radical innovation, and to a never-ending succession of radical and marginal innovation cycles. By contrast, without radical innovation, technology and knowledge converge to finite levels. Our next result shows that unlimited radical innovation never occurs. This result persists even when radical innovation entails an arbitrarily positive drift, as long as the optimal policy is well-defined, as shown in Section 6.

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.*

Thus, only marginal innovation persists along the equilibrium path. As the value of additional innovation goes to zero, we observe the emergence of a technological standard in the limit.[16]

A key step for understanding and proving Theorem 2 is to establish comparative statics for the incentives of radical innovation. Proposition 5 below shows that the *incremental* value and the size of radical innovation are both decreasing functions of z. The argument applies ideas from the experimentation literature to analyze our spatial model of technological space.

Let x^R denote the technology chosen when radical innovation takes place and $y_t^R(h_t) = x^R(h_t) - \bar{x}_t$ denote the optimal size of radical innovation.[17] Also let $\phi(\cdot)$ denote the density function of the standard normal distribution.

Proposition 5 (Radical Innovation) *If radical innovation is optimal, y_t^R solves*

$$\frac{\delta\sigma}{2\sqrt{y}}\phi\left(\frac{g_t}{\sigma\sqrt{y}}\right) = c'(y) \tag{5}$$

[16] This result bears some resemblance to the *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.

[17] There may exist several optima, although such case happens with zero probability. In such knife-edge cases, the comparative statics in the proposition still apply in the sense of the strong set order of lattice theory.

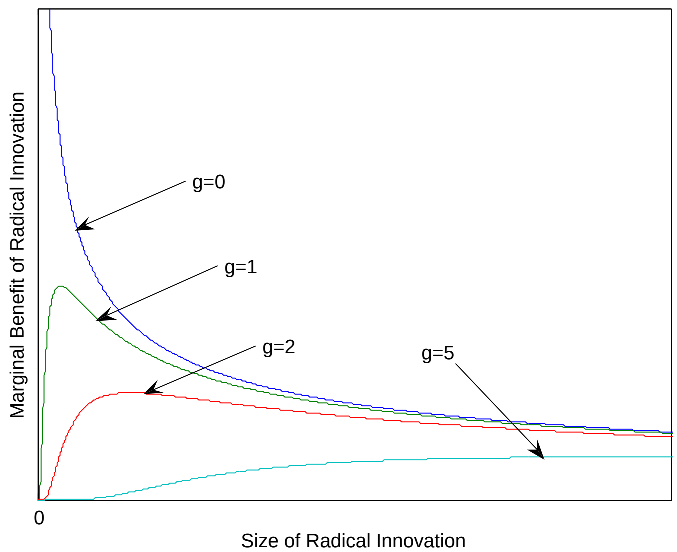

Figure 2: Marginal benefit of radical innovation.

and y_t^R is strictly increasing in σ and δ. If $g_t = 0$, then y_t^R is indeed strictly positive. If $g_t > 0$ and $y_t^R > 0$, then y_t^R is strictly decreasing in g_t. Finally, there exists a cutoff $\tilde{g} > 0$ above which $V^R(g) < 0$.

To get some intuition for Proposition 5, consider, first, the case of a zero gap. Radical innovations have an expected payoff of z_t, regardless of their size. A large size, however, comes with a higher variance and thus an increased probability of exceeding z_t. The incremental benefit of radical innovation (left-hand side of (5)) is arbitrarily large close to the frontier (see Figure 2). As the size of radical innovation increases, the volatility ($\sigma\sqrt{x - \bar{x}}$) increases at a decreasing rate, and the marginal benefit of radical innovation converges to zero. Since the marginal cost is increasing, the optimal size of radical innovation is therefore positive and well-defined.

The situation is quite different when the gap is positive, because 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 virtually 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. The incremental benefit of radical innovation is single peaked: it is initially pushed up by the increase in the probability of discovering a payoff above the current outside option z_t, which is given by $1 - \Phi\left(\frac{g_t}{\sigma\sqrt{x-\bar{x}}}\right)$, where Φ denotes the distribution function of a standard normal distribution. When the size of radical innovation reaches $\frac{g_t^2}{\sigma^2}$, the marginal benefit starts to decrease, as the probability of obtaining a payoff greater than z_t converges to $\frac{1}{2}$. Figure 2 illustrates the marginal benefit of radical innovation for different sizes of the gap.

Because volatility is the agent's only hope of improving the best technology, a higher volatility has a positive effect on the incentives to innovate. Similarly, a higher discount factor increases the incentives to innovate. By contrast, an increase in the gap reduces the incremental 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 hints at the reason 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. That proposition alone, however, does *not* imply that radical innovation terminates in finite time, because the technologies observed in equilibrium are chosen endogenously. To prove stagnation, a key step is to show that, if the value of radical innovation were positive at all times, the frontier would keep expanding in steps that are bounded below away from zero.[18]

Figure 3 shows a simulated path for the technology payoff function f, along with the equilibrium innovation dynamics for two cost specifications. The figure highlights several interesting features of our model. Innovation cycles arise endogenously along the equilibrium path. Both the length and the number of cycles is path dependent. In particular, it is not necessarily the case that a lower cost of radical innovation leads to more radical innovation. The figure shows that radical innovation may indeed end sooner when experimentation is cheaper. Cheaper experimentation leads to a larger size of radical innovation when the gap is zero but this larger size may lead an agent to explore an unattractive part of the technology space which is exactly what happens in the simulation. In general, a wave of radical innovations can go on for several periods before radical innovation ends forever. Eventually, the search for a better technology begins to cluster around a technology discovered by a previous gener-

[18] See Lemma 2. Intuitively, 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.

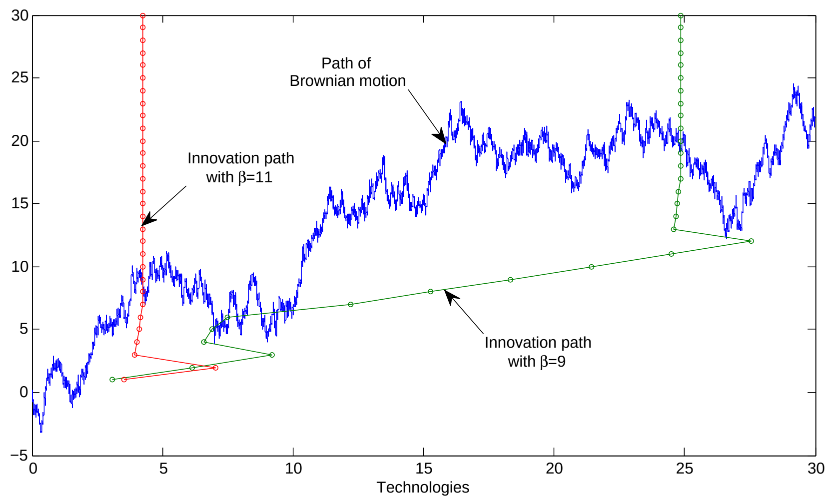

Figure 3: Dynamics of innovation: $\delta = 1$, $c(y) = \frac{y^2}{2\beta}$. The payoff function is the realized path of a Brownian motion with zero drift and volatility $\sigma = 3$.

ation, yielding the endogenous-information equivalent of an informational cascade. Society converges to a suboptimal technology, which is largely path dependent: a single experiment at, say, $x = 10$ would completely change the dynamics of the search process when radical innovation is cheaper, shifting innovation towards a different part of the technological space which would uncover better technologies.

5 Stagnation and Social Inefficiency

Our stagnation result points to a general inefficiency problem with social experimentation. Radical innovation generates two positive externalities on future generations, by increasing knowledge and by expanding the active domain. If these externalities were taken into account by current generations, wouldn't they systematically push radical innovation beyond the equilibrium level characterized in earlier sections? We formalize this question by introducing an infinitely-lived social planner who discounts payoffs with a discount factor $\delta_S < 1$.

Comparing the equilibrium dynamics of innovation with the social optimum raises several difficulties. Firstly, it is well-known that optimal experimentation with an infinitely lived agent and correlated technologies cannot be characterized by the arm-specific index policies that Gittins and Jones (1974) have identified for the standard multi-armed bandit problem with independent arms. This is intuitive: trying a new technology teaches something about surrounding technologies and affects the set of units.[19] Secondly, any comparison between equilibrium and efficient experimentation has to be of a probabilistic nature. Indeed, it is easy to construct specific Brownian paths for which equilibrium innovation will last longer than the social optimum, by assigning an implausibly low payoff to the first technology tried by the social planner and, repeatedly high payoffs for the radical innovations arising in equilibrium. The resulting histories are consistent with possible Brownian paths, and clearly result in equilibrium radical innovations lasting longer than what is prescribed by the efficient policy. Thirdly, it is not clear how the social planner should discount, if at all, future generations.

We circumvent all those difficulties by considering a limiting result as δ_S goes to 1, i.e., when the social planner becomes infinitely patient. Let $\bar{x}^{FB}(\delta_S)$ denote the frontier at which a social planner with discount factor δ_S stops radical innovation. Despite our inability to characterize the social optimum notwithstanding, we can prove the following result.

Theorem 3 *As δ_S goes to 1,* $\text{Prob}\left(\lim_{\delta_S \to 1} \bar{x}^{FB}(\delta_S) = +\infty\right) = 1$.

Theorem 3 implies, as a corollary, that equilibrium innovation is with high probability inefficiently low compared to the social optimum.

6 Stagnation with Optimistic Beliefs

Theorem 2 assumed that the drift κ which determines the expected payoff of new technologies, was equal to zero. What happens if agents are more optimistic about κ? Could that be enough to sustain radical innovation in the long run?

For this problem to have a well-defined solution, we must assume that the incremental cost of radical innovation eventually exceeds the drift κ. Otherwise, young agents may want to choose an infinite amount of radical innovation. Precisely, we assume that $\lim_{y \to +\infty} c'(y) > \kappa(1+\delta)$. For technical reasons, we also assume that $c''(0) > 0$.

[19] Readers may wonder whether units, which are conditionally independent, can be characterized by Gittins-like indices. We have explored this possibility and concluded that this was highly unlikely: an index of the kind that we obtained for two-period lived agents does not seem to exist for infinitely lived agents.

The scenario of a positive drift creates an incentive for old agents to innovate as well, as captured by the objective function

$$U^{O,R}(h_t) = \max_{x \in [\bar{x}_t, +\infty)} E_{h_t}\left[f(x) - c(x - \bar{x}_t)\right] = f(\bar{x}_t) + \kappa(x - \bar{x}_t) - c(x - \bar{x}_t) \qquad (6)$$

This slightly complicates the short-run equilibrium behavior of innovations, because agents now learn from both old and young generations. This does not, however, change the insight of Theorem 2.

THEOREM 4 *Radical innovation stops in finite time, almost surely.*

This extension of Theorem 2 is not completely straightforward, and may be explained as follows.

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

$$\kappa = c'(x - \bar{x}_t) \qquad (7)$$

Let $y^{O,R}$ denote the solution to (7) and $\xi = \kappa y^{O,R} - c(y^{O,R}) > 0$. An old agent prefers radical innovation over exploitation if and only if

$$f(\bar{x}_t) + \kappa y^{O,R} - c(y^{O,R}) \geq z_t, \text{ or, equivalently, if } g_t \leq \xi. \qquad (8)$$

When choosing the size of his innovation, a young agent considers the effect of his action today on his incentives tomorrow. This effect can be quantified in a perceived reduction of the gap from g_t to $g_t - \xi$ 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.

7 Discussion

Causes of Stagnation

Our stagnation theorem differs from the stagnation mechanisms already identified in the growth literature. For example, Aghion and Howitt (1992) identify a "no growth" equilibrium in which no research is carried out because the prospect of a high level of research in the future destroys current incentives for innovation, due to competition effects. Kortum (1998) develops a search-theoretic model in which successful innovations make the discovery of even better technologies harder to achieve. Our economy converges to stagnation without the requirement that better technologies are harder to discover as time goes by.[20]

Our analysis also captures Gordon's intuition that innovation dynamics are based on the interplay between radical innovations and their refinement through marginal innovation. Radical innovations eventually get stuck as future generations become unwilling to make the necessary investments to further push the technological frontier. However, it is the willingness of each generation to undertake bold innovations that fades away over time, not the possibilities to innovate radically. As the slowdown in productivity is the product of myopic behavior by self-interested agents, it leaves open the possibility that a more forward-looking agent, like a government, could actually help stimulate and sustain growth by either subsidizing or directly investing in radical innovation.

Our stagnation result continues to hold if marginal innovations are also costly. We do need to maintain the assumption that *exploitation* is costless. This natural assumption is, among others, justified by Romer (1990, p. 72): "once the cost of creating a new set of instructions has been incurred, the instructions can be used over and over again at no additional cost. Developing new and better instructions is equivalent to incurring a fixed cost. This property is taken to be the defining characteristic of technology." In the Appendix (Section E), we study the incentives for radical innovation when the cost c of radical innovation also depends, negatively, on the value of the best explored technology, which captures the idea that old technologies can also be used as inputs for discovering new ones. Our stagnation result continues to hold in that case.

There are many mechanisms affecting long-run innovation, which may precipitate or delay stagnation. Our stagnation result can most easily be circumvented by making the cost of radical innovation time dependent, either through exogenous shocks or deterministically as a result of population growth (as in Kortum (1998)), or by assuming, as in Jovanovic and Rob (1990), that there is a pool of new technologies to choose from, at any time, whose payoffs are independent from past innovations. More generally, one could assume that the expected benefit of new technologies depends on more than the payoff at the frontier (for example,

[20] See also Jones (1995), Segerstrom (1998), and Zeira (2011).

by introducing a multidimensional technological set). One interpretation of our theorem is that modeling features of this kind are necessary for radical innovation to be sustained in the long run. In general, our theorem shows that the intuition, from the experimentation literature, that forward-looking agents may fail to learn optimum actions, takes a particularly strong form in the context of radical innovation with an unbounded technological space, as stagnation occurs with certainty.

Patents and Radical Innovation

One way of improving inefficiently short radical innovation is to use patents. In our model, patents would consist of transfers between consecutive generations and would be defined by exploiting the spatial nature of the technological space. In particular, one could 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.

Such a patent protection system can affect incentives for radical innovation through two channels:

1. Getting Royalty Fees Radical innovation now brings a positive probability that the payoff 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, as shown by Proposition 4. When this happens, a small royalty fee just paid by the new generation increases the ex ante incentives for radical innovation. Perhaps counter-intuitively, though, even in the case of a flat royalty (let alone a more complex royalty structure), the level of that royalty affects the incremental value of radical innovation and, therefore, the optimal size of radical innovation. Moreover, this value may be *non-monotonic* in the patent level.[21]

2. 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

[21] 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.

is "pushed" towards further radical innovation.

These incentives are substitutes of each other: the more the new generation avoids royalty fees (say, because they are fixed at high levels), and the lower the patent incentives for the old generation, and vice versa. However, regardless of which effect dominates, royalty fees, structured in this way, would foster radical innovation.

Infinitely-Lived Firms

We modeled firms as two-period agents. This choice is guided by tractability: it is well-known and easy to show that, with correlated arms, optimal experimentation is impossible to characterize by simple indices, even for a single agent. Our two-period firms do capture, in a tractable fashion, the standard exploration/exploitation trade-off that is at the heart of experimentation. What we do not capture is the strategic considerations that infinitely lived firms would have regarding information acquisition, such as free riding and encouragement effects. Free riding might occur because a firm can avoid radical innovation in the hope that some other firm will incur the cost of expanding the current frontier. At the same time, however, radical innovation pursued by one firm may encourage other firms to pursue more radical innovation when previous experimentation led to the discovery of an attractive technology. We refer the reader to the work of Bolton and Harris (1999), Keller, Rady, and Cripps (2005), and Keller and Rady (2010) for key models along these lines.

Appendix

The Appendix is organized as follows. Appendix A contains the proofs of all the results stated in Sections 3 and 4, except for the proof of Theorem 2, which is contained in Appendix B. The proofs of Theorems 3 and 4 are contained in Appendix C and D, respectively. Appendix E extends the baseline model to the case in which the cost of radical innovation depends on the outcomes of past innovations. Finally, Appendix F shows that the baseline model can be interpreted as the reduced form of a competition game between multiple firms.

A Omitted Proofs

Proof of Theorem 1

Given that the distribution of payoffs within each unit is conditionally independent from other units, we first compute the value of innovation of a bounded unit u. Suppose, without loss of generality, that $M = f(x_r) > f(x_l) = m$ (the case of an equality is obtained by taking the limit as, say, m is increased to M). Then, for any $x \in [x_l, x_r]$,

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

where $\mathcal{N}(\kappa, \sigma^2)$ denotes the distribution of a Gaussian variable with mean κ and variance σ^2. Letting $a(x) = \frac{x - x_l}{x_r - x_l}$,

$$f(x) - m \sim \mathcal{N}\left(d\ a(x), a(x)(1 - a(x))L\sigma^2\right)$$

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

$$\begin{aligned} E\max\{k(x), z'\} = {} & z'\Phi\left(\frac{z' - da(x)}{\sigma\sqrt{a(x)(1 - a(x))L}}\right) + da(x)\left[1 - \Phi\left(\frac{z' - da(x)}{\sigma\sqrt{a(x)(1 - a(x))L}}\right)\right] \\ & + \sigma\sqrt{a(x)(1 - a(x))L}\phi\left(\frac{z' - da(x)}{\sigma\sqrt{a(x)(1 - a(x))L}}\right) \end{aligned}$$

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

$$
\begin{aligned}
E_u\left[f(x)+\delta \max\{f(x),z\}\right] &= (1+\delta)m + da(x) + \delta E\max\{k(x),z'\} \\
&= (1+\delta)m + d\left\{a(x)\left(1+\delta-\delta\Phi\left(\frac{\frac{z'}{d}-a(x)}{\sigma\frac{\sqrt{L}}{d}\sqrt{a(x)(1-a(x))}}\right)\right)\right. \\
&\quad + \delta\sigma\sqrt{a(x)(1-a(x))}\frac{\sqrt{L}}{d}\phi\left(\frac{\frac{z'}{d}-a(x)}{\sigma\frac{\sqrt{L}}{d}\sqrt{a(x)(1-a(x))}}\right) \\
&\quad \left. +\delta\frac{z'}{d}\Phi\left(\frac{\frac{z'}{d}-a(x)}{\sigma\frac{\sqrt{L}}{d}\sqrt{a(x)(1-a(x))}}\right)\right\} \qquad (9)
\end{aligned}
$$

Define $q_1 \equiv \frac{\sqrt{L}}{d}$, $q_2 \equiv \frac{z'}{d}$, and

$$
\begin{aligned}
\bar{\eta}\left(a,q_1,q_2\right) = \frac{1}{1+\delta}&\left\{a\left(1+\delta-\delta\Phi\left(\frac{q_2-a}{\sigma q_1\sqrt{a(1-a)}}\right)\right)\right. \\
&\left. +\delta q_2\Phi\left(\frac{q_2-a}{\sigma q_1\sqrt{a(1-a)}}\right)+\delta\sigma\sqrt{a(1-a)}q_1\phi\left(\frac{q_2-a}{\sigma q_1\sqrt{a(1-a)}}\right)\right\} \qquad (10)
\end{aligned}
$$

As the common argument of Φ and ϕ depends only on $a(x)$, $\sqrt{L}/d$ and z'/d, we can write

$$
E_u\left[f(x)+\delta\max\{f(x),z\}\right] = (1+\delta)\left\{m + d\,\bar{\eta}\left(a(x),\frac{\sqrt{L}}{d},\frac{z'}{d}\right)\right\} \qquad (11)
$$

Maximizing the expected payoff over $x \in [x_l, x_r]$ yields

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

where

$$
\eta\left(q_1,q_2\right) = \max_{a\in[0,1]}\bar{\eta}\left(a,q_1,q_2\right). \qquad (13)
$$

Letting $V(u,z) = \left[m + d\,\eta\left(\frac{\sqrt{L}}{d},\frac{z-m}{d}\right)\right] - z$ yields the characterization of the value of innovation for a bounded unit.

Next, notice that

$$
\begin{aligned}
\frac{\partial\bar{\eta}(a,q_1,q_2)}{\partial a} &= \frac{1}{1+\delta}\left\{1+\delta-\delta\Phi\left(\frac{q_2-a}{\sigma q_1\sqrt{a(1-a)}}\right)+\frac{\delta\sigma q_1}{2}\frac{1-2a}{\sqrt{a(1-a)}}\phi\left(\frac{q_2-a}{\sigma q_1\sqrt{a(1-a)}}\right)\right\} \\
&\geq \frac{1}{1+\delta} > 0
\end{aligned}
$$

for any $a \leq \frac{1}{2}$, because the sum of the first three terms in curly brackets is weakly greater than 1 by the property of Φ and the last term is positive if, and only if, $a \leq \frac{1}{2}$. Thus, any $a^* \in \arg\max_{a\in[0,1]} \bar{\eta}(a, q_1, q_2)$ must be greater than $\frac{1}{2}$. As $a(x) = \frac{x-x_l}{x_r-x_l}$, the corresponding optimal technology from a unit with interval $[x_l, x_r]$, $x^*(u, z)$, 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 some 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).[22] 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, that[23]

$$\frac{\partial V(u,z)}{\partial L} = \frac{\delta}{1+\delta}\frac{\sigma}{2}\sqrt{\frac{a^*(1-a^*)}{L}}\phi\left(\frac{z'-da^*}{\sigma\sqrt{a^*(1-a^*)L}}\right) > 0 \tag{14}$$

where $a^* \in \arg\max_{a\in[0,1]} \bar{\eta}\,(a, q_1, q_2)$. Similarly, for $M \neq z$, $\frac{\partial V(u,z)}{\partial M} > 0$ and $\frac{\partial V(u,z)}{\partial m} > 0$, as can be seen from equations (20) and (21) in the proof of Proposition 4.

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

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

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

$$E \max\{k(x), g\} = \sigma\sqrt{x-\bar{x}}\phi\left(\frac{g}{\sigma\sqrt{x-\bar{x}}}\right) + g\Phi\left(\frac{g}{\sigma\sqrt{x-\bar{x}}}\right). \tag{15}$$

[22] The 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).

[23] An 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.

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

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

Defining

$$\eta^\infty(g) = \frac{1}{1+\delta}\sup_{y\geq 0}\left\{ -c(y) + \delta\sigma\sqrt{y}\phi\left(\frac{g}{\sigma\sqrt{y}}\right) + \delta g\Phi\left(\frac{g}{\sigma\sqrt{y}}\right)\right\} \tag{16}$$

and $V(u,z) = [f(\bar{x}) + \eta^\infty(g)] - z = \eta^\infty(g) - g$ yields the formula claimed by Theorem 1. The value of innovation 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 value of innovation.■

Proof of Proposition 1

For any bounded unit u with $M < z$ we have[24]

$$\frac{d\eta(q_1,q_2)}{dz} = \frac{1}{d}\frac{\partial\eta(q_1,q_2)}{\partial q_2} = \frac{\delta}{(1+\delta)d}\Phi\left(\frac{q_2 - a^*}{\sigma q_1\sqrt{a^*(1-a^*)}}\right) \tag{17}$$

where, as before, $q_1 \equiv \frac{\sqrt{L}}{d}$, $q_2 \equiv \frac{z-m}{d}$, and $a^* \in \arg\max_{a\in[0,1]} \bar{\eta}(a, q_1, q_2)$. This implies that

$$\frac{\partial V(u,z)}{\partial z} = \frac{\delta}{1+\delta}\Phi\left(\frac{z-m-da^*}{\sigma\sqrt{a^*(1-a^*)L}}\right) - 1 < 0.$$

If, instead, u contains the best explored technology, then $V(u,z) = m - z + (z-m)\eta\left(\frac{\sqrt{L}}{z-m}, 1\right)$. This yields $\frac{\partial V(u,z)}{\partial z} = \eta\left(\frac{\sqrt{L}}{z-m}, 1\right) - \frac{\sqrt{L}}{z-m}\frac{\partial\eta}{\partial q_1}\left(\frac{\sqrt{L}}{z-m}, 1\right) - 1$, which is strictly negative from (22). 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,

$$V(u,z_t) < 0 \iff \frac{U(u;h_t)}{1+\delta} < z_t \iff U(u;h_t) < (1+\delta)z_t$$

where the second equivalence follows from (12). Therefore, choosing the best explored 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. ■

[24]Again, we omit dependence on side-derivatives. See Footnote 22 and the discussion surrounding it.

Proof of Proposition 2

Suppose that $M = z$. Differentiating $\bar{\eta}(a, q_1, 1)$ with respect to a yields

$$(1+\delta)\frac{\partial\bar{\eta}(a,q_1,1)}{\partial a} = \left[1+\delta-\delta\Phi\left(\frac{1}{\sigma q_1}\sqrt{\frac{1-a}{a}}\right)\right] + \frac{\delta\sigma q_1}{2}\frac{1-2a}{\sqrt{a(1-a)}}\phi\left(\frac{1}{\sigma q_1}\sqrt{\frac{1-a}{a}}\right)$$

which tends to $-\infty$ as a goes to 1. Also, $\bar{\eta}(0, q_1, 1) \equiv \lim_{a\to 0}\bar{\eta}(a, q_1, 1) = \frac{\delta}{1+\delta} < 1 = \lim_{a\to 1}\bar{\eta}(a, q_1, 1) \equiv \bar{\eta}(1, q_1, 1)$. We have thus shown that any solution to (13) lies in (0,1) whenever $q_1 > 0$ and $q_2 = 1$. Consider any history h_t and unit u with z as one of its endpoint payoffs. Then, $d = z - m$ and

$$V(u,z) = m-z+(z-m)\eta\left(\frac{\sqrt{L}}{z-m},1\right) > m-z+(z-m)\bar{\eta}\left(1,\frac{\sqrt{L}}{z-m},1\right) = m-z+(z-m) = 0$$

which proves that there is always at least one bounded unit with strictly positive value of innovation. Thus, exploitation is strictly suboptimal. ■

Proof of Proposition 3

We need to evaluate the effect of an increase in the size of the gap, g, on the value of radical innovation $V^R(g) = \eta^\infty(g) - g$. We suppose once again that the value function (16) is differentiable. If there exists an interior solution $y^* > 0$ to (16) when $g > 0$, then $\frac{\partial V^R(g)}{\partial g} = \frac{d\eta^\infty(g)}{dg} - 1 = \frac{\delta}{1+\delta}\Phi\left(\frac{g}{\sigma\sqrt{y^*}}\right) - 1 < 0$. If there is no interior solution, then $\eta^\infty(g) = \frac{\delta}{1+\delta}g$ which gives $\frac{\partial V^R(g)}{\partial g} = \frac{\delta}{1+\delta} - 1 < 0$. When $g = 0$, $V^R(g) = \eta^\infty(0)$ and $\frac{\partial V^R(g)}{\partial g} = 0$.

Finally, if $V^R(g_t) < 0$, 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 best explored technology. ■

Proof of Proposition 4

If $V^R(g_t) > V^M(h_t)$, then radical innovation occurs with probability 1 at time t. Thus, consider a history h_t such that $V^M(h_t) > V^R(g_t) > 0$. By continuity, there exists $\varepsilon > 0$ such that the value of radical innovation remains positive if the current gap is increased to $g_t + \varepsilon$. 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 $B(\epsilon)$.

Next, suppose that radical innovation occurs at an arbitrary history h_t. Let f_t denote the corresponding payoff of radical innovation, and let u_t denote the newly created bounded unit. For the rest of this proof, it turns out to be more convenient to work with the auxiliary indexes $\gamma(u,z) = V(u,z) + z = m + d\;\eta\,(q_1,q_2)$ and $\gamma(u^\infty,z) = \eta^\infty(g) + f(\bar{x})$ which are a simple normalization of the value function of a young agent when his first period choice is restricted to a unit u. Thus, the young agent experiments with a technology within the unit with the highest (auxiliary) index $\gamma(u,z)$. In particular, we can define the indexes for marginal and radical innovation by

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

for any history h_t, where $\mathcal{P}(h_t)$ is the collection of bounded units induced by a history h_t. After radical innovation has taken place, the new index for marginal innovation at the start of time $t+1$ is

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

We first consider how $\bar{\gamma}(\max\{z_t,f_t\})$ varies with f_t. If $f_t \le z_t$, $\bar{\gamma}(\cdot)$ is unaffected by f_t. If $f_t > z_t$, consider any unit $u \in \mathcal{P}(h_t)$. Then,[25]

$$\frac{\partial\gamma(u,f_t)}{\partial f_t} = \frac{\partial}{\partial f_t}\left[m + d\eta\left(\frac{\sqrt{L}}{d}, \frac{f_t - m}{d}\right)\right] = \frac{\partial \eta}{\partial q_2}\left(\frac{\sqrt{L}}{d}, \frac{f_t-m}{d}\right) \in \left(0, \frac{\delta}{1+\delta}\right) \tag{19}$$

by (17). Next, we consider the index of the new unit u_t. To this end, define f_l as the payoff associated with the left endpoint of unit u_t. If $f_l < f_t \le z_t$, then

$$\begin{aligned}\frac{\partial\gamma(u_t,z_t)}{\partial f_t} &= \frac{\partial}{\partial f_t}\left[f_l + (f_t - f_l)\eta\left(\frac{\sqrt{L}}{f_t-f_l}, \frac{z_t-f_l}{f_t-f_l}\right)\right] \\ &= \frac{a^*}{1+\delta}\left[1+\delta-\delta\Phi\left(\frac{q_2-a^*}{\sigma q_1\sqrt{a^*(1-a^*)}}\right)\right] > 0 \end{aligned} \tag{20}$$

[25]Again dropping dependence on side-derivatives, see Footnote 22.

where $a^* \in \arg\max_{a\in[0,1]} \bar{\eta}\,(a, q_1, q_2)$. If, instead, $f_t < \min\{z_t, f_l\}$, then

$$\begin{aligned}\frac{\partial\gamma(u_t, z_t)}{\partial f_t} &= \frac{\partial}{\partial f_t}\left[f_t + (f_l - f_t)\eta\left(\frac{\sqrt{L}}{f_l - f_t}, \frac{z_t - f_t}{f_l - f_t}\right)\right] \\ &= (1 - a^*)\left[1 - \frac{\delta}{1+\delta}\Phi\left(\frac{q_2 - a^*}{\sigma q_1\sqrt{a^*(1-a^*)}}\right)\right] \in \left(0, \frac{1}{1+\delta}\right) \end{aligned} \tag{21}$$

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

$$(1 - a^*)\left[1 - \frac{\delta}{1+\delta}\Phi\left(\frac{q_2 - a^*}{\sigma q_1\sqrt{a^*(1-a^*)}}\right)\right] \leq \frac{1}{2}\left[1 - \frac{\delta}{1+\delta}\frac{1}{2}\right] = \frac{2+\delta}{4(1+\delta)} < \frac{1}{1+\delta}$$

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

$$\frac{\partial\gamma(u_t, f_t)}{\partial f_t} = \frac{1}{1+\delta}\left[a^*(1+\delta) + \delta(1-a^*)\Phi\left(\frac{1}{\sigma q_1}\sqrt{\frac{1-a^*}{a^*}}\right)\right] \in \left(\frac{\delta}{1+\delta}, 1\right) \tag{22}$$

The upper bound follows from $a^* < 1$, since $q_2 = 1$ in this case, while the lower bound follows from $a^* \geq \frac{1}{2}$, by Theorem 1.

Below z_t, $\bar{\gamma}$ is flat as a function of f_t, while the index of the new unit strictly increases in f_t, by (20) and (21). Above z_t, $\bar{\gamma}$ grows at most by $\frac{\delta}{1+\delta}$ following an increase in f_t, while the derivative of the index of the new unit is at least $\frac{\delta}{1+\delta}$ as shown by (22). If $\bar{\gamma}(z_t) > \gamma(u_t, z_t)|_{f_t=z_t}$, the index of the new unit is strictly below $\bar{\gamma}(z_t)$ for any $f_t < z_t$. 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 $\bar{\gamma}(z_t) < \gamma(u_t, z_t)|_{f_t=z_t}$, the two indexes necessarily cross only once at some $f_t < z_t$, but never above z_t. We ignore the case $\bar{\gamma}(z_t) = \gamma(u_t, z_t)|_{f_t=z_t}$, which occurs with zero probability. The intersection is denoted by f_t^M and it is such that marginal innovation would occur, when optimal, in the newly created unit if and only if $f_t > f_t^M$.

We now analyze how the index of radical innovation varies with f_t, which again denotes the new frontier payoff. With a slight abuse of notation, the index of radical innovation at time $t+1$ is given by $\gamma_{t+1}^R(f_t) = f_t + \eta^\infty(g_{t+1}) = f_t + \eta^\infty\left(\max\{(z_t - f_t), 0\}\right)$ as a function of the frontier payoff f_t. For any $f_t \geq z_t$, $\gamma_{t+1}^R(f_t) = f_t + \eta^\infty(0)$ and $\frac{\partial\gamma_{t+1}^R}{\partial f_t} = 1$.

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

For $f_t \in (\tilde{f}, z_t)$, the maximizer of (16) may be interior. In that case, $\frac{\partial \gamma^R_{t+1}}{\partial f_t} = 1 - \frac{\delta}{1+\delta}\Phi\left(\frac{z_t - f_t}{\sigma\sqrt{y^*}}\right) \in \left(\frac{1}{1+\delta}, 1\right)$, 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: $g_t = 0$. Let f^M_t denote the cutoff such that marginal innovation occurs in the new unit if, and only if, $f_t \geq f^M_t$, whose existence we just proved. If $f^M_t \geq z_t$, the index of marginal innovation is flat for any $f_t < z_t$, and

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

which follows from our previous analysis. Since the slope of the index of radical innovation is always at least $\frac{1}{1+\delta} \geq \frac{\delta}{1+\delta}$ below f^M_t and 1 above f^M_t (as $f^M_t \geq z_t$), it follows that γ^R_{t+1} and γ^M_{t+1} as functions of f_t cross exactly once. Let $f^R_t = f^R(h_t)$ denote such intersection.

Similarly, if $f^M_t < z_t$, the index for radical innovation is unchanged, but

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

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

Case 2: $g_t > 0$. If $f^M_t \geq z_t$, the analysis is the same as for Case 1. Thus, there exists a unique intersection f^R_t.

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

Proof of Proposition 5

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

$$E_h\left[f(\bar{x}+y) - c(y) + \delta \max\{f(\bar{x}+y), z\}\right] = (1+\delta)f(\bar{x}) - c(y) + \delta\sigma\sqrt{y}\phi(0),$$

from (16). The first-order condition yields Equation (5) with $g = 0$. The right-hand side of (5) 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 $y \to +\infty$. Therefore, there always exists a solution to Equation (5) when $g = 0$, 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 $g > 0$, the expected utility from radical innovation of size $y > 0$ is

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

The first-order condition is again given by (5). Differentiating (5) with respect to g, we obtain $-\frac{\delta g}{2\sigma y^{3/2}}\phi\left(\frac{g}{\sigma\sqrt{y}}\right)$, which is strictly negative. Since the objective function in (16) is submodular in (y, g), the Strict Monotonicity Theorem 1 of Edlin and Shannon (1998) implies that the optimal size of radical innovation is decreasing in g. Similarly, the objective function is supermodular in (y, δ) and (y, σ).

We now show the existence of a cutoff $\tilde{g}$ above which radical innovation has a negative value. Let $A(y, g) \equiv \frac{\delta\sigma}{2\sqrt{y}}\phi\left(\frac{g}{\sigma\sqrt{y}}\right)$ denote the marginal benefit of radical innovation, given a gap g and a size y of radical innovation.[26]

Lemma 1 *For any $g > 0$,*

1. $\lim_{y\to+\infty} A(y, g) = \lim_{g\to+\infty} A(y, g) = \lim_{y\to0} A(y, g) = \lim_{y\to0} \frac{\partial A(y,g)}{\partial y} = 0$;
2. $\frac{\partial A(y,g)}{\partial g} < 0$.
3. *$A(y, g)$ is strictly quasi-concave in y, and maximized at $\frac{g^2}{\sigma^2}$.*

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 $\lim_{y\to0} A(y, g)$ is computed using the fact that $\lim_{y\to0} \frac{1}{\sqrt{y}}e^{-\frac{1}{y}} = \lim_{z\to\infty} \frac{z^{1/2}}{e^z} = 0$. One shows similarly that $\lim_{y\to0} \frac{\partial A(y,g)}{\partial y} = 0$. Strict quasiconcavity of A in y comes from the fact that

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

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

[26] $A(y, g)$ corresponds to the left-hand side of equation (5).

To establish the existence of the threshold $\tilde{g} > 0$, we start by showing that $y^* = 0$ is the unique maximizer of (16), whenever g is high enough. If $c'(0) > 0$, Lemma 1 implies that $\lim_{y \to 0} A(y, g) = 0 < c'(0)$ for any $g > 0$. Also by Lemma 1, $\frac{\partial A(y,g)}{\partial g} < 0$, $\lim_{g \to +\infty} A(y, g) = 0$, and $A(y, g) \leq A\left(\frac{g^2}{\sigma^2}, g\right) = \frac{\delta\sigma^2}{2g}\phi(1)$. This, combined with the properties of c, implies the existence of a large enough threshold $\tilde{g}$ such that $c'(y) > A(y, g)$ for all $g > \tilde{g}$ and $y > 0$. A similar argument applies if $c'(0) = 0$ and $c''(0) > 0$, because $\lim_{y \to 0} \frac{\partial A(y,g)}{\partial y} = 0$, from Lemma 1. Substituting $y^* = 0$ into (16) yields $\eta^\infty(g) = \frac{\delta g}{1+\delta}$. Therefore, the value of radical innovation is equal to $V^R(g) = -\frac{g}{1+\delta} < 0$, for $g > 0$. ■

B 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 $\tilde{g} > 0$ such that radical innovation ends after any history such that $g_t \geq \tilde{g}$. 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, $y^R_{\varphi(t)} = x^R_{\varphi(t)} - \bar{x}_{\varphi(t)} > 0$ 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 $y^R_{\varphi(t)}$ converges to 0. From Proposition 5, $y^R_{\varphi(t)}$ is decreasing in the gap. Therefore, there must exist a subsequence $\{\psi(t)\}$ of $\{\varphi(t)\}$ such that $\{g_{\psi(t)} = z_{\psi(t)} - f(\bar{x}_{\psi(t)})\}$ is increasing. Since that sequence is bounded above by $\tilde{g}$, it must converge to some strictly positive limit ρ, and $z_{\psi(t)}$ converges to the limit $\rho + f(\tilde{x})$. For sufficiently high t, the expected payoff from radical innovation is approximately equal to $f(\tilde{x}) + \delta\tilde{z}$, while the payoff 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 best explored technology yielding a payoff close to $\tilde{z}$ 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

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

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

Proof. By the recurrence property of Brownian motion, there exists a.s. an $\tilde{x} > 0$ such that $f(\tilde{x}) < -\tilde{g}$ and, hence, some $\zeta > 0$ such that $A_\zeta(-\tilde{g}) > 0$. 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 $-\tilde{g}$. Because z is always nonnegative, the gap $z - f(\bar{x})$ after such history will exceed $\tilde{g}$, 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. 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) $\lim_{x\to-\infty} G(x, \Lambda) = -\Lambda$ and $\lim_{x\to+\infty} G(x, \Lambda) = +\Lambda$. Such function is easily built, and is bounded by Λ.

For any Brownian path f, the transformed path $b^\Lambda : x \mapsto G(f(x), \Lambda)$ is continuous and bounded by Λ, and is homeomorphic to f. In particular, b^Λ and f are observationally equivalent. Any belief μ about f has a corresponding belief μ^Λ about b^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 b^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 $Z^{\Lambda}(\mu) = \sup\{b^{\Lambda}(x_t) : x_t \text{ contained in } h\}$. 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 $Z(\mu) = \max\{f(x_t)\} = G^{-1}(Z^{\Lambda}(\mu), \Lambda)$[27] 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 payoff 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

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

We will use the following lemma, which is proved at the end of this Appendix (Section B.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.

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

Proposition 6 *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 $f(x) = Z(\mu)$.*

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

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

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 $x, \hat{\mu}, z$

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

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 by[28]

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

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.[29] We will establish a stronger result, whose proof is in Appendix B.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 $\bar{Z}$.

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

[29]The 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.

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

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

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 $Z(\mu_t)$.[30] 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 6 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.

B.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 $(x, \mu, z) \in$

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

$K \times \Delta(\tilde{\Theta}) \times [-\Lambda, \Lambda]$. Then,

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

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

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

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

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

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

$$
A\left(\frac{1}{j}, \varepsilon\right) = \left\{ b^\Lambda \in \tilde{\Theta} :\ \exists\, \lambda > \frac{1}{j} \text{ s.t. } |x - y| < \lambda \Longrightarrow |b^\Lambda(x) - b^\Lambda(y)| < \varepsilon \right\}
$$

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

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

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

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

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

B.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 $\Lambda \to \infty$, 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.[31] To show uniform integrability, it suffices to prove that there exists $p > 1$ such that:[32]

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

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

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

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

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

[32] See, e.g., Durrett (1996), Exercise 4.5.1., p. 260.

is the maximum of a Brownian bridge with width δ_i (the width of π) and endpoints equal to 0. This implies that[33]

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

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

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

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

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

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

C Proof of Theorem 3

We first need to define the values of marginal and radical innovation for the social planner. Given a discount factor δ_S and an arbitrary history h_t, define the following auxiliary functions:[34]

$$\begin{aligned} U^{FB,M}(h_t; \delta_S) \equiv \quad & \sup_{\{x_r\}_{r=t}^{+\infty}} E_{h_t} \left[f(x_t) + \sum_{r=t+1}^{+\infty} \delta_S^{r-t} \left(f(x_r) - c(x_r - \bar{x}_r) \right) \right] \\ s.t. \quad & x_t \in [0, \bar{x}_t], \; x_r \in \mathbb{R}_+, r = t+1, ... \end{aligned} \tag{25}$$

[33]See, 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.

[34]Recall the convention that $c(y) = 0$ for any $y \leq 0$.

and

$$U^{FB,R}(h_t;\delta_S) \equiv \sup_{\{x_r\}_{r=t}^{+\infty}} E_{h_t}\left[f(x_t) - c(x_t - \bar{x}_t) + \sum_{r=t+1}^{+\infty} \delta_S^{r-t}\left(f(x_r) - c(x_r - \bar{x}_r)\right)\right] \quad (26)$$
$$s.t. \quad x_t \geq \bar{x}_t,\ x_r \in \mathbb{R}_+, r = t+1, ...$$

$U^{FB,M}(h_t;\delta_S)$ is the value function of a social planner with discount factor δ_S when his time-t choice is restricted to technologies that lie within the frontier associated with an arbitrary history h_t. Similarly, $U^{FB,R}(h_t;\delta_S)$ represents the value function when the time-t choice must be greater than or equal to the frontier. As for the baseline model, if the social planner finds optimal to exploit the best explored technology at an arbitrary time t, then he will stick to that technology forever. This is because no learning occurs from exploitation so that the social planner would start $t+1$ with exactly the same beliefs that made exploitation optimal in the previous period. Thus, despite the higher complexity of the social planner's problem, we can define the values of marginal and radical innovation in a similar way to the baseline model up to normalization:

$$V^{FB,M}(h_t;\delta_S) \equiv U^{FB,M}(h_t;\delta_S) - \frac{z_t}{1-\delta_S}, \quad \text{and} \quad V^{FB,R}(h_t;\delta_S) \equiv U^{FB,R}(h_t;\delta_S) - \frac{z_t}{1-\delta_S} \quad (27)$$

As the values of marginal and radical innovation are simply an affine transformation of the value functions of the social planner's maximization problem (restricting the domains of his first choice within and outside the frontier, respectively), it also follows that the social planner prefers marginal over radical innovation at a history h_t if, and only if, $V^{FB,M}(h_t;\delta_S) \geq V^{FB,R}(h_t;\delta_S)$.

We start the analysis with a technical result.

Lemma 6 *There exists a sequence $\{\tilde{g}(\delta_S)\}$ with $\tilde{g}(\delta_S) \to +\infty$ as $\delta_S \to 1$ such that the value of radical innovation for a social planner with discount factor δ_S is positive for any gap $g \leq \tilde{g}(\delta_S)$.*

Proof. Let g and $\bar{x}$ denote the gap and the frontier associated with an arbitrary history h. Define a 1-step policy as a policy such that the social planner experiments only once and then exploits forever after. Given a history h, let $V^{1,R}(g;\delta_S)$ denote the value of radical innovation when the social planner is restricted to using only a 1-step policy at any history

h with associated gap g.[35] Clearly, $V^{FB,R}(h;\delta_S) \geq V^{1,R}(g;\delta_S)$ as a 1-step policy is a feasible policy for the social planner. The restriction to 1-step policies has the advantage that we can explicitly write down the associated maximization problem as

$$
\begin{aligned}
V^{1,R}(g;\delta_S) &= \sup_{y\geq 0}\left\{\frac{f(\bar{x})}{1-\delta_S} - c(y) + \frac{\delta_S}{1-\delta_S}\sigma\sqrt{y}\phi\left(\frac{g}{\sigma\sqrt{y}}\right) + \frac{\delta_S}{1-\delta_S}g\Phi\left(\frac{g}{\sigma\sqrt{y}}\right)\right\} - \frac{z}{1-\delta_S} \\
&= \frac{1}{1-\delta_S}\sup_{y\geq 0}\left\{-g - (1-\delta_S)c(y) + \delta_S\sigma\sqrt{y}\phi\left(\frac{g}{\sigma\sqrt{y}}\right) + \delta_S g\Phi\left(\frac{g}{\sigma\sqrt{y}}\right)\right\} \qquad (28)
\end{aligned}
$$

which is equivalent to the maximization problem for an agent in our baseline model with a discount factor equal to $\frac{\delta_S}{1-\delta_S}$. Note that, for $\delta_S = 1$, the maximization problem reduces to,[36]

$$
\sup_{y\geq 0}\left\{-g + \sigma\sqrt{y}\phi\left(\frac{g}{\sigma\sqrt{y}}\right) + g\Phi\left(\frac{g}{\sigma\sqrt{y}}\right)\right\} = +\infty \qquad (29)
$$

Thus, $V^{1,R}(g;1) = +\infty$, for any $g \geq 0$. This implies that there exists a $\bar{\delta} \in (0,1)$ such that, for any $\delta_S \geq \bar{\delta}$, we can find a positive number $\tilde{g}(\delta_S)$ such that $V^{1,R}(\tilde{g}(\delta_S);\delta_S) > 0$. By Proposition 3, it also follows that $V^{FB,R}(h;\delta_S) \geq V^{1,R}(g;\delta_S) > 0$, for any $g \leq \tilde{g}(\delta_S)$.

Suppose by way of contradiction that there is no sequence $\{\tilde{g}(\delta_S)\}$ with $\delta_S > \bar{\delta}$ as defined in the proof of Lemma 6 such that the sequence diverges. Take any such sequence. Then, there must exist a $\bar{g} > 0$ such that $\tilde{g}(\delta_S) < \bar{g}$ along that sequence.[37] This implies that there also exists $g' > \bar{g}$ and a subsequence such that $V^{1,R}(g';\delta_S) < 0$ for any δ_S along the subsequence. However, evaluating (29) at $g = g'$ implies that $V^{1,R}(g';\delta_S)$ will eventually be positive as $\delta_S \to 1$. We thus reached a contradiction. ■

Fix a path of Brownian motion. Let $\bar{x}^{FB}(\delta_S)$ denote the frontier at which a social planner with discount factor δ_S will stop radical innovation. Suppose by way of contradiction that $\liminf_{\delta_S\to 1}\bar{x}^{FB}(\delta_S) = \hat{x} < +\infty$. Take a subsequence $\bar{x}^{FB}(\delta_{S,n})$ such that $\bar{x}^{FB}(\delta_{S,n}) \uparrow \hat{x}$ and $\tilde{g}(\delta_{S,n}) \to +\infty$ as $n \to +\infty$.[38] For $\hat{x} < +\infty$ to be the case, it must be that the sup norm of the path of Brownian motion over the interval $[0, \bar{x}^{FB}(\delta_{S,n})]$, defined as $\sup_{x\in[0,\bar{x}^{FB}(\delta_{S,n})]}|f(x)|$, exceeds $\frac{\tilde{g}(\delta_{S,n})}{2}$ for any n. Suppose to the contrary that there exists $\bar{n}$ such that $\sup_{x\in[0,\bar{x}^{FB}(\delta_{S,\bar{n}})]}|f(x)| \leq \frac{\tilde{g}(\delta_{S,\bar{n}})}{2}$. By Lemma 6, the value of radical innovation for a

[35] When restricted to 1-step policies, the value of radical innovation depends on the history only through the gap as in the baseline model. This is not the case for general policies which can instead depend on the history in a non-trivial way.

[36] The derivative of the objective function with respect to y is $\frac{\sigma}{2\sqrt{y}}\phi\left(\frac{g}{\sigma\sqrt{y}}\right) > 0$.

[37] If necessary, we can restrict attention to a subsequence.

[38] Such a subsequence exists by Lemma 6.

social planner with discount factor $\delta_{S,\bar{n}}$ would be bounded away from zero over $[0, \bar{x}^{FB}(\delta_{S,\bar{n}})]$, regardless of the history, because the associated gap can never exceed $\tilde{g}(\delta_{S,\bar{n}})$ over that interval. As the social planner keeps experimenting within the fixed interval $[0, \bar{x}^{FB}(\delta_{S,\bar{n}})]$, he will eventually explore every unit that still generates a positive value of innovation.[39] This is because the frontier stays the same during experimentation and the payoff function over $[0, \bar{x}^{FB}(\delta_{S,\bar{n}})]$ is bounded with probability 1 by the almost sure continuity of Brownian motion, as in the baseline model. Therefore, the value of marginal innovation for the social planner must eventually converge to zero and, consequently, below the value of radical innovation. This means that the social planner will also eventually find radical innovation optimal, which leads to a contradiction. Hence, it must be the case that $\sup_{x \in [0, \bar{x}^{FB}(\delta_{S,n})]} |f(x)| > \frac{\tilde{g}(\delta_{S,n})}{2}$ for every n along the subsequence that we are considering.

Finally, recall that $\tilde{g}(\delta_{S,n}) \to +\infty$ as $n \to +\infty$ (or $\delta_{S,n} \to 1$). This implies that the path of Brownian motion must have an infinite sup norm over the interval $[0, \bar{x}]$ but this cannot be the case because the a.s. continuity of Brownian motion implies that $\sup_{x \in [0,\bar{x}]} |f(x)| < +\infty$ with probability 1. We thus reached the desired contradiction. Finally, we can conclude that $\liminf_{\delta_S \to 1} \bar{x}^{FB}(\delta_S) = +\infty$ with probability 1. ■

D Proof of Theorem 4

We start by showing the existence of a threshold $\tilde{g}$ above which $y^R(g) = 0$. Suppose, first, that $\kappa \leq c'(0)$, so that an old agent chooses the best explored technology (as observed in the main text). The expected utility of a young agent from choosing technology $x > \bar{x}$ is

$$f(\bar{x}) + \kappa y - c(y) + \delta \left\{ (f(\bar{x}) + \kappa y) \left(1 - \Phi\left(\frac{g - \kappa y}{\sigma \sqrt{y}} \right) \right) + \sigma \sqrt{y} \phi \left(\frac{g - \kappa y}{\sigma \sqrt{y}} \right) \right. \quad (30)$$

$$\left. + (g + f(\bar{x}))\, \Phi\left(\frac{g - \kappa y}{\sigma \sqrt{y}} \right) \right\} \quad (31)$$

$$= (1 + \delta)[f(\bar{x}) + \kappa y] - c(y) + \delta \left\{ (g - \kappa y) \Phi\left(\frac{g - \kappa y}{\sigma \sqrt{y}} \right) + \sigma \sqrt{y} \phi \left(\frac{g - \kappa y}{\sigma \sqrt{y}} \right) \right\}, \quad (32)$$

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

$$\kappa \left(1 + \delta - \delta \Phi\left(\frac{g - \kappa y}{\sigma \sqrt{y}} \right) \right) + \frac{\delta \sigma}{2\sqrt{y}} \phi \left(\frac{g - \kappa y}{\sigma \sqrt{y}} \right) = c'(y). \quad (33)$$

[39] Even if we have not explicitly defined the value of innovation of a unit, such a concept can be easily defined following (25) and (27).

The left-hand side of (33) approaches κ as g increases and, for fixed g, the left-hand side converges to $\kappa(1+\delta)$ as $y \to +\infty$, and to κ as $y \to 0$. Also, the left-hand side is bounded above by

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

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

Differentiating (33) with respect to g yields

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

for all $g, y > 0$. This implies that $y^R(0) \geq y^R(g)$ for all $g > 0$. 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 $g \leq \xi$, as shown in the main text. We assume without loss of generality that $g > \xi$, so that an old agent does not innovate today and then the expected utility today of any $x > \bar{x}$ for a young agent is simply

$$E[f(x) - c(x-\bar{x})] + \delta\{E[f(x)+\xi|f(x) \geq z-\xi]\ \mathrm{Prob}(f(x) \geq z-\xi) + z\ \mathrm{Prob}(f(x) < z-\xi)\}$$
$$= (1+\delta)(f(\bar{x}) + \kappa y) - c(y) + \delta\left\{\sigma\sqrt{y}\phi\left(\frac{g-\xi-\kappa y}{\sigma\sqrt{y}}\right) + \xi + (g-\xi-\kappa y)\Phi\left(\frac{g-\xi-\kappa y}{\sigma\sqrt{y}}\right)\right\}$$

The first-order condition is

$$\kappa\left[1+\delta-\delta\Phi\left(\frac{g-\xi-\kappa y}{\sigma\sqrt{y}}\right)\right] + \frac{\delta\sigma}{2\sqrt{y}}\phi\left(\frac{g-\xi-\kappa y}{\sigma\sqrt{y}}\right) = c'(y) \tag{34}$$

The right-hand side is always greater than or equal to κ. Since $g_t - \xi > 0$, the right-hand side converges to κ as $y \to 0$, and to $\kappa(1+\delta)$ as $y \to +\infty$. Since also $c'(0) < \kappa$, there must exist an interior solution to the first-order equation (34). When $\lim_{y\to+\infty} c'(y) > \kappa(1+\delta)$, the solution is unique for high values of g because the left-hand side of (34) approaches κ pointwise as g becomes arbitrarily large. As g increases, $y^R(g)$ approaches $y^{O,R}$: 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

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

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

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 $g > \xi$, the right-hand side of (34) is strictly decreasing in g. Thus, the unique positive solution of the first-order condition when $g = \xi$ provides the desired upper bound. When $g \leq \xi$, 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 $g \in [0, \xi]$, the value and size of radical innovation are decreasing in g, providing an upper bound on the size of radical innovation for a young agent, uniform over all histories.

E Cost Externalities

In the baseline model, the path of innovations affects incentives for radical innovation through i) the expected value of new technologies, and ii) the opportunity cost of forgoing marginal innovation. It is perhaps more realistic to allow radical innovation to be directly affected by past technologies, if those technologies are helpful as inputs to produce radical innovations. Our model may be extended along this direction, 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 w, the function $c(\cdot, w)$ satisfies the same assumptions as in the baseline model.

Let $y^R(g_t, \alpha z_t)$ denote the optimal size of radical innovation at a history h_t with associated gap g_t, and current best payoff z_t. The following result is established similarly to Proposition 5 and is stated without proof.[40]

Proposition 7 (Comparative statics) *Suppose that $\alpha > 0$ and that c is submodular. Then, $y^R(g, \alpha z)$ 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 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 $\bar{c}(y) = \lim_{z \to +\infty} c(y, \alpha z)$ denote the lower envelope of the cost functions $\{c(\cdot, w)\}$.[41]

Proposition 8 (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.*

Proof. Since $\bar{c}(\cdot)$ is increasing in a right neighborhood of $y = 0$, the properties of each cost function $c(\cdot, \alpha z)$ then guarantee that $\bar{c}(\cdot)$ is increasing everywhere. Replicating the proof of Proposition 5, one can show the existence of a threshold $\tilde{g} > 0$ above which the marginal

[40]The agent's objective function is submodular in (y, α). When there are multiple maximizers, Proposition 7 holds in the sense of the strong set order (see footnote 17).

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

benefit of radical innovation is strictly less than the marginal cost, at any $y > 0$, for any $z > 0$. Thus, $y^R(g, \alpha z) = 0$ for any $g > \tilde{g}$, 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 $g > 0$, 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 $\lim_{z \to +\infty} y^R(0, \alpha z) < \Lambda$. Finally, we can now repeat the same proof as in Step 1 of Theorem 2. ■

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.

F Competition Among Firms

Suppose that in each period there is a new cohort of consumers of unit mass and N firms that produce a good. Firms live for two periods and discount profits at a common rate $\delta \in (0, 1]$. New generations of firms *do not* overlap. Firms sell their good to consumers.[42]

Each firm n can experiment with a technology $x_n \in \mathbb{R}_+$ to produce the good which has a net (of input costs) productivity given by $f(x_n)$, so that $f(\cdot)$ can take on negative values. Each firm shares the same beliefs about the mapping between technologies and net productivity, $f(\cdot)$, which is represented by the realized path of a Brownian motion with drift $\kappa = 0$ and variance parameter σ^2. Experimentation with technologies that are in the active domain is costless because those technologies are combinations of known technologies. Each firm n instead incurs an experimentation cost $\frac{c(x_n - \bar{x})}{N}$ from experimenting radically with a technology $x_n > \bar{x}$.

Consumers have a unit demand and need the good. The firm that preempts all other firms in the development, production, and marketing of a good captures the entire market. In particular, given any vector of technology choices, $\mathbf{x} = (x_1, ..., x_N)$, we assume that each firm has an equal chance of preemption. If firm n wins the race, its profits gross of any experimentation costs are represented by $f(x_n)$. To further simplify the analysis, we assume

[42] If generations of firms where overlapping, it would create the problem that old firms would not be exploiting anymore because young entrepreneurial firms could produce a higher quality product that could reduce their market share to zero.

that at most one firm wins the race in any period and that this firm sells to all consumers. Firms may also exploit known technologies. Even if the productivity of a technology that has been used in the past is known, a firm that has never used that technology may face problems in its implementation. We capture this possibility by assuming that a firm that employs a known technology will also face an equal chance of preemption as any other firm. Firms can freely imitate any known technology but the productivity of a new technology discovered by a young firm is observable only to that firm until the new cohort of firms is born.

This formulation of competition among firms follows Loury (1979). Loury assumes a constant reward for the first firm that introduces an innovation and firms can choose the size of R&D which determines the likelihood of being the first innovator. We, instead, assume that each firm has an equal probability to preempt its rivals but the "size" of the reward is stochastic. In particular, a negative $f(x)$ may be interpreted as the case in which a firm overspent in the development of the product compared to the consumers' perceived quality of that product. A firm may still prefer to sell to the entire market in order to recover some fixed costs.

The next result shows that our baseline model can be interpreted as a reduced-form version of a more complex model in which several firms compete with each other.

Proposition 9 *Competition among firms generates the same equilibrium path as our baseline model with overlapping generations of single agents.*

Proof. Suppose that the industry is composed of old firms. Each old firm solves, with the convention that $c(y) = 0$ for $y \leq 0$,

$$\max_{x_n \geq 0} \; E_h \left[\frac{f(x_n)}{N} - \frac{c(x_n - \bar{x})}{N} \right] \tag{35}$$

As each old firm has an equal probability of being the first to develop an innovation or market the good using a known technology, each firm only cares about developing the product which guarantees the highest expected profit upon winning the race. Thus, each old firm will exploit the best explored technology as in the baseline model.

Let $\mathbf{x}_{-n}$ denote the vector of technologies chosen by young firms $j \neq n$. Given that each old firm behaves as in the baseline model, we can write the maximization problem of a young firm as

$$\max_{x_n \geq 0} \; E_h \left[\frac{f(x_n)}{N} - \frac{c(x_n - \bar{x})}{N} + \delta \frac{\max\{f(x_n), z\}}{N} \right] \tag{36}$$

The continuation profits follow from the assumption that the young firm can observe the productivity of any technology it explored, and those explored by other firms except for the technologies discovered by firms in its own cohort.

The maximization problem (36) is equivalent to (3) up to normalization. Whereas there could be multiple optimal technologies following an arbitrary history h, given (36) it is straightforward to show that for any such solution we can construct a symmetric equilibrium in which each firm experiments with that technology.

References

Aghion, P., and P. Howitt, 1992, "A Model of Growth Through Creative Destruction," *Econometrica*, 60, 323–351.

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.

Bolton, P., and C. Harris, 1999, "Strategic Experimentation," *Econometrica*, 67, 349–374.

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.

Gittins, J., and D. Jones, 1974, "A Dynamic Allocation Index for the Sequential Allocation of Experiments," in *Progress in Statistics*, ed. by J. Gani. North-Holland, Amsterdam, pp. 241–266.

Gordon, R., 2012, "Is U.S. Economic Growth Over? Faltering Innovation Confronts the Six Headwinds," *NBER working paper series*, Working paper 18315.

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 Institue, Washington, DC, 140–83.

Jones, C., 1995, "R&D-Based Models of Economic Growth," *Journal of Political Economy*, 103, 759–784.

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.

Keller, G., and S. Rady, 2010, "Strategic Experimentation with Poisson Bandits," *Theoretical Economics*, 5, 275–311.

Keller, G., S. Rady, and M. Cripps, 2005, "Strategic Experimentation with Exponential Bandits," *Econometrica*, 73, 39–68.

Kortum, S., 1998, "Research, Patenting, and Technological Change," *Econometrica*, 65, 1389–1419.

Lee, I., 1993, "On the Convergence of Informational Cascades," *Journal of Economic Theory*, 61, 395–411.

Loury, G. C., 1979, "Market Structure and Innovation," *Quarterly Journal of Economics*, 93, 395–410.

Milgrom, P., and I. Segal, 2002, "Envelope Theorems for Arbitrary Choice Sets," *Econometrica*, 2, 583–601.

Romer, P., 1990, "Endogenous Technological Change," *Journal of Political Economy*, 98, 71–102.

Segerstrom, P., 1998, "Endogenous Growth without Scale Effects," *Amreican Economic Review*, 88, 1290–1310.

Slywotsky, A., 2009, "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," *BusinessWeek*, August 27th.

Zakaria, F., 2011, "The Future of Innovation: Can America Keep Pace?," *NewsWeek*, June 5th.

Zeira, J., 2011, "Innovations, Patent Races and Endogenous Growth," *Journal of Economic Growth*, 16, 135–156.