Northwestern University
2001 Sheridan Road 580 Leverone Hall Evanston, IL 60208-2014 USA

Discussion Paper #1516
December 9, 2010

"The Dynamics of Innovation"

Key words: innovation, Schumpeter cycles, experimentation, social learning, R&D, intergenerational externalities

JEL classification: D83, D92, 03, 04

Umberto Garfagnini
Northwestern University

Bruno Strulovici
Northwestern University

www.kellogg.northwestern.edu/research/math

CMS-EMS
The Center for
Mathematical Studies
in Economics &
Management Sciences

The Dynamics of Innovation*

– Preliminary –

Umberto Garfagnini
Kellogg School of Management

Bruno Strulovici
Northwestern University

December 9, 2010

Abstract

We analyze social learning and innovation in an overlapping generations model in which available technologies have correlated payoffs. Each generation experiments within a set of policies whose payoffs are initially unknown and drawn from the path of a Brownian motion with drift. Marginal innovation consists in choosing a technology within the convex hull of policies already explored and entails no direct cost. Radical innovation consists in experimenting beyond the frontier of that interval, and entails a cost that increases with the distance from the frontier, and may decrease with the best technology currently available. We study how successive generations alternate between radical and marginal innovation, in a pattern reminiscent of Schumpeterian cycles. Even when the underlying Brownian motion has a positive drift, radical innovation stops in finite time. New generations then fine-tune policies in search of a local optimum, converging to a single technology. Our analysis thus suggests that sustaining radical innovation in the long-run requires external intervention.

*Email addresses: u-garfagnini@kellogg.northwestern.edu b-strulovici@northwestern.edu.

1 Introduction

The invention of the transistor in the 1950s, the first large-scale cellular-phone system in the 1970s, and the internet in the 1990s are all major historical breakthroughs that boosted the development of new industries for at least the next decade. The importance of innovation from a social perspective has ample empirical evidence (Griliches (1992), Hall (1996)), but is such innovation sustainable in the long-run?

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

We analyze a model of social learning and innovation where available policies (or "technologies") have correlated payoffs. At each period, a new, "young" generation can experiment with a new technology chosen on the positive real half-line, whose payoff is initially unknown and drawn from the path of a Brownian motion with drift.[2] Each generation lives for two periods and thus has an incentive to experiment in its first period. We distinguish between radical and marginal innovation: *marginal innovation* consists in choosing a technology that is in the convex hull (i.e., interval) of policies that have been tried earlier. Marginal innovation bears a fixed cost, which is normalized to zero. *Radical innovation* consists in experimenting beyond the frontier of that interval, and incurs a convex cost that increases with the distance from the frontier, and may decrease with the best technology currently available. We study how successive generations alternate between radical and marginal in-

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

[2]The usual time dimension of Brownian motion corresponds to the domain of alternatives which each generation can choose from.

novation, in a pattern reminiscent of Schumpeterian cycles.

The main finding of the paper shows that radical innovation stops in finite time with probability one, even when the underlying Brownian motion has a positive drift, and may be explained as follows. New discoveries increase the available opportunities to experiment marginally. If innovation has been very successful in the past, an agent may be more inclined to pursue a marginal rather than a radical innovation. The comparative advantage of a marginal innovation lies in a more predictable outcome at no cost. Similarly, the discovery of a bad technology at the frontier might shift attention towards marginal innovation. The result is robust even when the underlying Brownian motion has a positive drift. Thus, the threat to long-run growth follows either from the failure of the innovation process or from its own success.

Once the "stagnation" stage has been reached, new generations merely fine-tune policies in search of a local optimum. We show that marginal innovation always dominates exploitation, that is, the adoption of technologies already chosen by previous generations. Exploitation is always suboptimal in the short-run, because the refinement of the best known technology guarantees a higher expected payoff. However, the "value" of marginal innovation converges to zero in the limit, as the opportunities for meaningful innovation become increasingly rare. Limit behavior converges to a single technology, thus exploitation becomes optimal only in the long-run.

The paper proposes a formal definition for the *value of marginal innovation* and the *value of radical innovation*. Since an agent can always free ride on the experiences of his predecessors, the value of each type of innovation is represented by the expected benefit in excess of the best available technology, which evolves endogenously over time. The value of innovation drives incentives in the short-run, and determines which type of innovation is chosen in each period. Using the value of innovation, we are able to provide a precise description for the dynamics of innovation, which fully takes into account the fact that future innovation depends on current technologies. A comparative statics argument shows that the value of *marginal* innovation always increases in the best available technology, while the value of *radical* innovation may actually decrease following the discovery of a better technology. Indeed, discovering a better technology sets a higher bar on the expected outcome of radical innovation.

We disentangle the various intergenerational linkages occurring with social learning. First, there is a *direct learning effect* that arises from the observability of the outcomes of innovation by previous generations. Innovation increases the stock of knowledge of society through more refined beliefs about the underlying outcome process. Thus, direct learning has a positive effect on the well-being of future generations. Second, a *feasibility effect* arises when a new radical innovation is undertaken, because it endows the society with new opportunities for costless marginal innovation. A third intergenerational linkage arises when the cost of radical innovation is reduced following new discoveries. This *cost effect* creates additional incentives for radical innovation.

Our paper contributes to the literature on neo-Schumpeterian growth models. We build on the seminal contribution by Jovanovic and Rob (1990), who studied a search-theoretic model of growth through technological innovation. We depart from Jovanovic and Rob (1990) by endogenizing the dependence between incentives for future innovation and current technologies. In our model, the incentives to perform radical innovation are highly sensitive to the history of previous discoveries in two ways. First, there is a payoff externality that derives from the possibility of exploitation of previous technologies, as in Jovanovic and Rob (1990). Second, the history affects directly the beliefs that each generation holds about the outcome of radical innovation, which is completely absent in Jovanovic and Rob (1990).[3] In a different framework, Matsuyama (1999) investigates the connection between neoclassical and neo-Schumpeterian growth models, while in a more recent paper Acemoglu, Gancia, and Zilibotti (2010) study the interplay between innovation and standardization in a dynamic general equilibrium model. By contrast, we analyze innovation and standardization (radical vs marginal innovation) in a framework affected by uncertainty about the payoff associated with a new technology.

The paper can also be interpreted as a theory of economic stagnation. Jovanovic and Nyarko (1996) show that the development of human capital specific to a particular technology may induce a (long-lived) agent to stick to that technology in the long-run, despite the availability of (possibly) better technologies. Garicano and Rossi-Hansberg (2009) provide (among other things) an alternative theory of stagnation, which is based on the necessity to develop orga-

[3]This aspect also relates our model with the literature on organizational behavior, which stresses the tendency of firms to perform local searches around the technologies currently in use. This is a natural response to uncertainty about the outcomes of innovation. For example, Kauffman, Lobo, and Macready (2000) look at environments in which the firm's current location in the space of technologies influences the incentives for innovation.

nizations to exploit newly discovered technologies. In the present paper, stagnation follows endogenously from the outcomes of the innovation process.[4]

We contribute to the literature on optimal experimentation following Rothschild (1974), McLennan (1984), Easley and Kiefer (1988) and Aghion, Bolton, Harris, and Jullien (1991), among others. Some of our arguments build on Easley and Kiefer (1988) to derive the properties of the experimentation policy in the long run. In our setting, not only does there exist a well-defined long-run belief about the value of each policy: we show that experimentation converges to a single policy.[5] The literature has focused on the possibility of incomplete learning in the long-run and has tried to identify the conditions under which learning can be guaranteed to be complete. By contrast, learning can hardly be expected to be complete in our set-up, as the underlying parameter is the realized path of a Brownian motion. Our focus is on the effect of social learning on long-run behavior, once radical innovation has come to an end. In this regard, our paper is closer to Bala and Goyal (1998), but in an environment with a continuum of possible policies.[6]

2 Model

In each period $t = 0, 1, 2, ...$, a young decision maker is born, who lives for two periods. Each decision maker is risk-neutral and can experiment with a technology $x \in X = [0, \infty)$ in each period. Technologies are mapped into monetary outcomes according to the continuous function $f : X \to \mathbb{R}$, which is initially unknown except for the default option $f(0) = 0$.[7]

A possible interpretation of the space X is as the set of technologies available to produce a single good. A decision maker learns about the profitability of a new technology only

[4]Callander (2009) studies a model of dynamic policymaking in a similar framework to ours and shows that society might settle with undesirable policies.

[5]Our setting allows us to prove the convergence result without the compactness assumptions made in Easley and Kiefer (1988).

[6]The assumption that the outcome of past decisions is publicly known sets our model apart from the standard social learning literature by Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992). While in those models full observability would render the problem trivial, in our framework the knowledge of even arbitrarily many outcomes does not convey enough information to recover the underlying outcome function.

[7]The analysis can easily be extended to the entire real line.

through innovation, which could either improve on previous technologies or be less profitable. In either case, choices are *irreversible*: the decision maker commits himself to the particular technology he chooses for that period.[8]

The history $h_t = \{(0,0), (x_0, f(x_0)), \left((x_1^O, f(x_1^O)), (x_1^Y, f(x_1^Y))\right), ..., \left((x_{t-1}^O, f(x_{t-1}^O)), (x_{t-1}^Y, f(x_{t-1}^Y))\right)\}$ contains all the information about the technologies chosen by previous generations, young and old, until time t. Histories are publicly observable. We let $\mathcal{H}_t$ denotes the set of histories up to time t.

Each history induces a partition of the technology space X. Following Strulovici (2010), we use the following definition.

Definition 1 *A unit* $u = [I, m, M]$ *is comprised of an interval* $I = [x_l, x_r]$, *where* $m = \min\{f(x_l), f(x_r)\}$, *and* $M = \max\{f(x_l), f(x_r)\}$. *Also,* $l = x_r - x_l$ *and* $d = M - m$.

If $l = +\infty$, *then a unit* $u_\infty = [I, m]$ *is simply defined as an interval* $I = [x_l, +\infty)$, *and* $m = f(x_l)$.

A unit can be interpreted as a subproblem in its own right. We denote by $\mathcal{P}(h_t)$ the collection of finite units induced by an arbitrary history h_t. The next definition contains basic nomenclature used throughout the paper.

Given any history h_t, we introduce the following notation.

- $z_t = \max\{f(x_{t'}^O), f(x_{t'}^Y) : t' \leq t\}$ is the value of the technology explored by time t;
- $\hat{x}_t$ is the technology associated with z_t;[9]
- $x_t^f = \max\{x_{t'}^O, x_{t'}^Y : t' \leq t\}$ is the right boundary of the set of explored technologies, and will be referred to as the *frontier*;
- $\rho_t = z_t - f(x_t^f)$ is called the *gap*.

z is weakly increasing over time. A young decision maker can always obtain an expected payoff of at least z in each period. The frontier x^f is the upper bound on the feasible set

[8]This distinguishes our model from Jovanovic and Rob (1990), where the decision maker unveils the value of a new technology, and then decides whether or not to use it.

[9]That technology is unique with probability one.

of technological possibilities. While the gap measures how far behind the frontier stands compared to the state-of-the-art technology.

Each generation shares the common belief that the underlying outcome function is the realized path of a Brownian motion with drift $\mu = 0$ and volatility parameter $\sigma > 0$.[10] The Brownian structure of the problem has specific implications for the outcome distribution of unknown technologies. On the one hand, every action $x > x^f$ generates an outcome

$$f(x) \sim N\left(f(x^f), \sigma^2(x - x^f)\right) \tag{1}$$

Note that the current history h_t influences the distribution only through the frontier. On the other hand, an untried technology x lying in a unit u of finite lenght is normally distributed as well, but with mean

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

and variance

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

where x_l and x_r are the endpoints of the unit. This distribution is called a Brownian bridge.[11] Note how the choice of a technology affects only the distribution of technologies lying in the same unit. Also, the initial drift does not play any role.

This paper focuses on two types of innovation. First, *radical innovation* refers to the choice of technologies to the right of the frontier. Radical innovation can be interpreted as basic research, which helps expanding the space of feasible technologies. Second, *marginal innovation* involves technologies inside a unit within the frontier, and it corresponds to the improvement of old technologies. In addition, a decision maker could decide to *exploit* a previously tried technology, which guarantees a sure payoff.

Both exploitation and marginal innovation are costless, as we are assuming a convex feasibility condition on the space of technologies. That is, whenever the frontier expands, all intermediate technologies become feasible. Radical innovation incurs a cost $c(x - x^f)$, which depends on how far a decision maker pushes research away from the frontier. Intuitively, the cost might be interpreted as a necessary investment to back up basic research. We maintain the following assumption throughout the paper:

[10] The assumption that $\mu = 0$ makes the analysis neater but it could be dispensed with. See Section 7.

[11] See Billingsley (1968) for an extensive treatment.

Assumption 1 *$c(\cdot)$ is twice continuously differentiable, strictly increasing, weakly convex, and such that $c(0) = 0$.*

Section 6 analyzes the case where the cost of radical innovation decreases in the best technology currently available.

Each decision maker maximizes the expected present value of his payoffs in both periods. The discount factor $\delta \in (0, 1)$ is the same across generations. An old decision maker has no value for information and thus never undertakes radical innovation. The description of a history may thus be more simply expressed as $h_t = \{(0, 0), (x_0, f(x_0)), ..., (x_{t-1}, f(x_{t-1}))\}$, with the interpretation that $(x_{t'}, f(x_{t'}))$ is the technology-outcome pair chosen by the young generation at time $0 \leq t' < t$.

Finally, we can write the maximization problem of a young decision maker as follows,

$$V(h_t) = \sup_{x \in X} E_{h_t} \left[f(x) - c(x - x_t^f) + \delta \max\{f(x), z_t\} \right], \quad \forall\, h_t \in \mathcal{H}_t \tag{4}$$

with the convention that $c(y) = 0$ whenever $y \leq 0$. Since an old decision maker plays no active role in the benchmark model, we will refer to a young decision maker simply as a decision maker.

3 The Value of Innovation

Each generation faces the choice of a technology out of a continuum of possible alternatives. Each technology produces a deterministic outcome which is correlated to the outcomes of technologies that are close by, because of the continuity of the underlying outcome function.

Lemma 1 provides a more useful way to think about the decision problem faced by each decision maker. Technologies are bundled together according to the unit they belong to. Since the choice of a technology only changes the distribution of technologies that reside in the same unit, the decision maker faces the reduced problem of picking one out of finitely many units. Lemma 1 shows that we can characterize each unit based on a single index. The index of a unit u represents the largest expected payoff attainable from technologies lying within the unit.[12]

Lemma 1 (Index Characterization) *Fix a history $h_t \in \mathcal{H}_t$. A young decision maker assigns an index $\gamma(u,z)$ to each unit induced by h_t. In particular, there exist functions $\eta : \mathbb{R}^2 \to \mathbb{R}$, and $\psi : \mathbb{R}_+ \to \mathbb{R}$ such that*

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

The dependence of the index on the best technology is related to the possibility of free-riding, which comes from exploitation. Since the decision maker knows that he will exploit in the next period, his choice today determines his payoff today but also his second period payoff. Also, tomorrow's payoff changes if and only if the outcome of today's innovation exceeds z. For a finite unit, the ratio $\frac{z-m}{d} \geq 1$ measures the relative attractiveness of the unit compared to exploitation. When the ratio is 1, the unit has z as an endpoint outcome and therefore a marginal innovation is very appealing. If the ratio is very high, than innovation becomes less attractive because the chance of improving tomorrow's payoff is low. The ratio $\frac{\sqrt{l}}{d}$, instead, is specific to the unit and it is correlated with the slope of the Brownian bridge, which affects the expected payoff of every technology inside the unit.[13]

For an infinite unit, the gap is an important determinant of the incentives for radical innovation. Intuitively, a zero gap means that the distribution of outcomes to the right of the frontier has the mean centered on the current best technology z. Thus, incentives for radical innovation are at their highest. A positive gap, instead, implies that the current best technology lies strictly within the frontier, and the distribution of outcomes is then determined by the frontier outcome alone. Thus, radical innovation offers a lower expected payoff today compared to exploitation. However, radical innovation is still valuable because it offers a higher variance. The higher variance increases the likelihood that the realized outcome exceeds the current best technology, which increases tomorrow's expected payoff.

We can now use the index characterization to define the *value of innovation* of a unit u as

[12] The formulas for the indexes correspond to equations (18) and (16) in Appendix A.

[13] The square root has to do with the standard deviation inside a Brownian bridge, as can be seen from (3).

the difference between the unit's index and the current best technology. Formally,

$$K(u, z) = \gamma(u, z) - z \tag{5}$$

The value of innovation of a unit is the largest gain the decision maker can obtain from a unit in excess of the payoff guaranteed by exploitation. It is also convenient to formally define the indexes for marginal and radical innovation as follows,

$$\gamma^M(h_t) = \max_{u \in \mathcal{P}(h_t)} \gamma(u, z), \quad \text{and} \quad \gamma^R(h_t) = \gamma(u_\infty, z) \tag{6}$$

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

Definition 2 *The value of marginal (resp., radical) innovation is defined by* $K^M(h_t) = \gamma^M(h_t) - z_t$ *(resp.,* $K^R(h_t) = \gamma^R(h_t) - z_t$*).*

The value of marginal innovation is the maximal net expected maximal benefit from the improvement of old technologies. Likewise, the value of radical innovation is the net expected maximal benefit that can be obtained by experimentation with technologies to the right of the frontier.

Lemma 2 shows that the trade-off between marginal and radical innovation can be represented in a parsimonious way through the value of innovation. Thus, the dynamics of innovation can be analyzed *as if* there are only two alternatives in each period: marginal innovation and radical innovation.

Lemma 2 *Fix a history* h_t*. A decision maker prefers radical over marginal innovation if and only if* $K^R(h_t) > K^M(h_t)$*. Also, the value of marginal innovation is increasing in* z*, while the value of radical innovation is constant in* z *if* $\rho = 0$ *and decreasing in* z *if* $\rho > 0$*. Finally, if* $K^R(h_t) < 0$*, then radical innovation ends forever.*

An increase in z increases the value of marginal innovation. The value of innovation of a specific unit, however, may actually decrease, but there is always at least one finite unit that benefits from a higher outside option. For example, consider a unit u with $\frac{z-m}{M-m} > 1$ and suppose that the best technology improves from z to $z' > z$. The increase in z has no effect on the distribution of technologies inside the unit but it does affect the overall expected payoff. This latter effect makes innovation in the first period less appealing because now it has to produce an outcome above z' rather than z to be of any use in the second period. Thus, $K(u, z) < K(u, z')$. If, instead, $\frac{z-m}{M-m} = 1$, an increase in z affects the underlying distribution of all technologies in u by increasing the slope of the associated Brownian bridge, while leaving the variance unchanged. As a result, the value of innovation of that unit increases.

The value of radical innovation is completely determined by the gap, $K^R(h_t) = \psi(\rho_t) - \rho_t$. An increase in z leads to a wider gap which, in turn, depresses incentives to perform radical innovation. The reason for the lower incentives is twofold. First, an increase in z makes exploitation relatively more appealing. Second, radical innovation has become riskier because it is now required to produce a higher outcome than before to be of any value in the second period. Thus, if, at any time, the value of radical innovation becomes negative, radical innovation will always be dominated by exploitation thereafter.

We are now ready to state the main result of the paper.

THEOREM 1 *Radical innovation ends in finite time with probability one. After radical innovation has ended, the economy witnesses an infinite sequence of marginal innovations, which converge to a single technology. Thus, exploitation only occurs in the limit.*

The next section describes how incentives for either type of innovation vary in the short-run.

4 Short-Run Dynamics of Innovation

At time 0, the decision maker faces a trade-off between the initial default option $z = 0$ and radical innovation. Radical innovation is costly but it may lead to a higher payoff. While exploitation is costless but it constrains the agent's payoff tomorrow to 0. Proposition 1 shows that it is indeed optimal to innovate.

PROPOSITION 1 *At $t = 0$, the decision maker pursues a radical innovation and the optimal size of the innovation x_0^* uniquely solves*

$$\frac{\delta\sigma}{2\sqrt{2\pi x_0^*}} = \frac{dc(x_0^*)}{dy} \tag{7}$$

The size of the innovation is increasing in both σ and δ.

The intuition goes as follows. If the decision maker exploits in the first period, his discounted expected payoff is 0. Suppose, instead, that he decides to pursue radical innovation. Because of the zero drift condition, any action $x > 0$ has an expected payoff of 0, but it improves the probability of a higher payoff tomorrow from zero to $\frac{1}{2}$. Thus, the marginal benefit of innovation (the left-hand side of (7)) is very large close to the frontier, which makes exploitation unattractive. As the size of the radical innovation increases, the volatility increases at a decreasing rate (recall $\sigma\sqrt{x - x^f}$), and the marginal benefit of innovation converges to zero. Since the marginal cost is increasing, the optimal radical innovation is well-defined.

The size of the innovation turns out to be sensitive to the volatility of the underlying outcome process, because the volatility is the decision maker's only way to improve his expected payoff tomorrow. Thus, a higher volatility has a positive effect on the incentives to innovate. The same result holds true as the discount rate increases.

In the following period, a new generation is born. Since we assumed perfect observability of actions and outcomes, the new decision maker faces three available options: *i*) exploitation with $z_1 = \max\{f(x_0^*), 0\}$; *ii*) marginal innovation over $[0, x_0^*)$; or *iii*) radical innovation over $(x_0^*, +\infty)$. For expositional simplicity, suppose $f(x_0^*) > 0$ and $x_0^* = 1$. Thus, the new outside option is $z_1 = f(x_0^*)$. Also, (2) and (3) imply that every $x \in (0, x_0^*)$ has distribution $N(z_1 x, x(1-x)\sigma^2)$, and discounted expected utility

$$\begin{aligned} EU(x) &= z_1 x + \delta E[\max\{f(x), z_1\}] \\ &= (1+\delta) z_1 x + \delta z_1 (1-x)\Phi\left(\frac{z_1}{\sigma}\sqrt{\frac{1-x}{x}}\right) + \delta\sigma\sqrt{x(1-x)}\phi\left(\frac{z_1}{\sigma}\sqrt{\frac{1-x}{x}}\right) \end{aligned} \tag{8}$$

from the properties of the truncated normal distribution, where $\Phi(\cdot)$ and $\phi(\cdot)$ denote the CDF and pdf of the standard normal distribution. Differentiating with respect to x, we obtain

$$\frac{dEU(x)}{dx} = (1+\delta) z_1 - \delta z_1 \Phi\left(\frac{z_1}{\sigma}\sqrt{\frac{1-x}{x}}\right) + \delta\sigma \frac{1-2x}{2\sqrt{x(1-x)}}\phi\left(\frac{z_1}{\sigma}\sqrt{\frac{1-x}{x}}\right) \tag{9}$$

The last term of (9) converges to $-\infty$ as x approaches 1, which shows that it is optimal to set $x < 1$. Since $x = 1$ is feasible and guarantees $(1+\delta)z_1$, innovating marginally leads to a strictly higher payoff. Notice that the same conclusion holds even if $f(x_0^*) < 0$, but with $x = 0$ as a reference point.

Proposition 2 shows that the optimality of marginal innovation over exploitation holds more generally.

Proposition 2 *For any history $h_t \in \mathcal{H}_t$, the value of marginal innovation is always strictly positive.*

Furthermore, for any finite unit u with $f(x_l) < f(x_r)$, the optimal technology $x^M(u,z)$ lies in $\left(\frac{x_r+x_l}{2}, x_r\right)$, for any $z \geq f(x_r)$.

The decision maker displays a bias for technologies that are closer to the endpoint with the highest outcome. The reason is quite intuitive and it has to do with the properties of the Brownian bridge. Recalling (2) and (3), the expected value is increasing over the bridge in the direction of the more profitable endpoint, while the variance is concave and symmetric around the midpoint. Therefore, every technology below $\frac{x_r+x_l}{2}$ is strictly dominated by the corresponding technology in the upper part of the unit.

Does the time-1 decision maker prefer marginal or radical innovation? Obviously, the answer depends on the realized outcome $f = f(x_0^*)$. If f is very high, we can foresee that radical innovation is the optimal choice, while low realizations push towards marginal innovation. The indexes for the two options are, respectively,

$$\gamma^M(h_1) = \begin{cases} f\eta\left(\frac{\sqrt{x_0^*}}{f}, 1\right) & \text{if } f > 0 \\ \frac{\delta\sigma\sqrt{x_0^*}}{2\sqrt{2\pi}(1+\delta)} & \text{if } f = 0 \\ f - f\eta\left(\frac{\sqrt{x_0^*}}{-f}, 1\right) & \text{if } f < 0 \end{cases}$$

and

$$\gamma^R(h_1) = f + \psi(f^-),$$

where $f^- = \max\{-f, 0\}$. The index of radical innovation embeds the cost function in it, thus it is not necessarily the case that a positive realization makes radical innovation optimal. So, suppose first that $f > 0$, then $\frac{\partial \gamma^R}{\partial f} = 1$ while $\frac{\partial \gamma^M}{\partial f} \in \left(\frac{\delta}{1+\delta}, 1\right)$.[14] Therefore, the value of the indexes at $f = 0$ determines if there is going to be an intersection or not above $f = 0$. When $\psi(0) > \frac{\delta\sigma\sqrt{x_0^*}}{2\sqrt{2\pi}(1+\delta)}$, then radical innovation dominates marginal innovation for any $f \geq 0$. Otherwise, there exists $\hat{f} > 0$ such that radical innovation is optimal if and only if

[14] The latter is shown in the proof of Lemma 3.

$f > \hat{f}$.[15] Next, we consider $f < 0$. It can be shown that $\frac{\partial \gamma^R}{\partial f} \in \left[\frac{1}{1+\delta}, 1\right)$ while $\frac{\partial \gamma^M}{\partial f} \in \left(0, \frac{1}{1+\delta}\right)$. Therefore, there is going to be an intersection if and only if $\psi(0) \leq \frac{\delta\sigma\sqrt{x_0^*}}{2\sqrt{2\pi}(1+\delta)}$. We can collect the previous observations in the following remark.

Remark There exists $\hat{f}$ such that the time-1 decision maker pursues radical innovation if and only if $f(x_0^*) > \hat{f}$.

After an arbitrary history h_t, the analysis is more involved. In particular, two possible scenarios may arise following a history h_t: *i)* $\rho_t = 0$; or *ii)* $\rho_t > 0$. In the first case, the distribution of outcomes to the right of the frontier is directly affected by z_t, while it is independent of z_t in the second case.[16] Thus, incentives differ in the two scenarios.

Proposition 3 *Suppose $\rho_t = 0$, then the value of radical innovation is strictly positive. Also, the optimal size of the innovation, $y_t^R = x_t^R - x_t^f$, is constant, and it uniquely solves*

$$\frac{\delta\sigma}{2\sqrt{2\pi y_t^R}} = \frac{dc(y_t^R)}{dy} \tag{10}$$

Next, suppose $\rho_t > 0$. If an interior solution, $y^R > 0$, exists, it solves

$$\frac{\delta\sigma}{2\sqrt{y^R}}\phi\left(\frac{\rho_t}{\sigma\sqrt{y^R}}\right) = \frac{dc(y^R)}{dy} \tag{11}$$

When $\rho_t = 0$, the decision maker faces exactly the same incentives for radical innovation as the first generation at time 0. The marginal benefit of innovation is very high close to the frontier because the probability of improving on the best technology jumps to $\frac{1}{2}$.

The situation is very different with a positive gap. Even with a positive gap, the marginal benefit of radical innovation approaches zero for very large radical innovations, but now also for innovations close the frontier. Close to the frontier, a radical innovation is accompanied by an inadequately low increase in the volatility, with almost no impact on the expected payoff tomorrow. For any given ρ, the marginal benefit is bell-shaped. It is initially pushed up by the increase in the probability of discovering an outcome above z_t, which is given by $1 - \Phi\left(\frac{\rho_t}{\sigma\sqrt{x - x^f}}\right)$. As the size of the innovation keeps increasing, though, the marginal benefit decreases.

[15] The indifference happening at $f = 0$ is a zero probability event, and it can be ignored without loss of generality.

[16] Recall (1).

Equation (11) also shows that an increase in ρ reduces the marginal benefit of radical innovation for any size of the innovation. Thus, a decision maker would have to choose technologies that are increasingly distant from the frontier, as ρ widens, but this becomes more and more expensive. A simple monotone comparative statics argument shows that an increase in the gap effectively reduces the optimal size of the innovation the decision maker is willing to undertake.

Proposition 4 (Comparative Statics) *Let $y^R(h_t) = x^R(h_t) - x_t^f$ denote the optimal size of the radical innovation following history h_t. If $y^R(h_t) > 0$, then $\frac{\partial y^R}{\partial \rho} < 0$.*[17]

How do the indexes for marginal and radical innovation change following a round of radical innovation? Radical innovation at time t changes the incentives for marginal innovation at time $t+1$ in two ways: *i*) it creates an additional finite unit; *ii*) it potentially increases the outside option. Lemma 3 characterizes the evolution of the index of marginal innovation which follows a history-dependent cutoff rule. If the outcome of radical innovation at time t exceeds the relevant cutoff, then the subsequent marginal innovation is chosen within the new unit.

Lemma 3 *Suppose $K^R(h_t) > K^M(h_t)$, for some history h_t. Then, there exists $\bar{f} = \bar{f}(h_t)$ such that*

$$\gamma^M(h_{t+1}) = \begin{cases} \max_{u \in \mathcal{P}(h_t)} \gamma(u, \max\{z_t, f\}) & \text{if } f \leq \bar{f} \\ \gamma(u'(f), \max\{z, f\}) & \text{if } f > \bar{f} \end{cases}$$

where $u'(f)$ is the new finite unit generated from radical innovation at time t with resulting outcome f.

Thus, high realizations shift attention towards the newly created unit. However, the new unit need not have the current best technology at time $t+1$ as an endpoint outcome. The reason is that the length of a unit is important too in determining its value of innovation, as can be seen from the index characterization in Lemma 1.

Radical innovation at time t also affects the new index for radical innovation. A sufficiently high outcome, for example above the current z_t, is good news for the time-$t+1$ decision

[17] When the first-order condition admits multiple solutions, the proposition can be stated in the strong set order sense.

maker and we might expect radical innovation to be optimal again. To the contrary, a very low realization makes marginal innovation look much more appealing. Once again, the gap plays a crucial role in our analysis.

Proposition 5 *Suppose $K^R(h_t) > K^M(h_t)$, for some history h_t.*

(i) If $\rho_t = 0$, then there exists a unique $\hat{f} = \hat{f}(h_t)$ such that $K^R(h_{t+1}) > K^M(h_{t+1})$ if and only if $f > \hat{f}$.

(ii) If $\rho_t > 0$, then we must distinguish between two cases:

(a) If $\bar{f}(h_t) \geq z_t$, where $\bar{f}(h_t)$ is given by Lemma 3, then there exists a unique $\hat{f} = \hat{f}(h_t)$ such that $K^R(h_{t+1}) > K^M(h_{t+1})$ if and only if $f > \hat{f}$.

(b) If $\bar{f}(h_t) < z_t$, then there exists $\hat{f} = \hat{f}(h_t)$ such that $f > \hat{f}$ implies $K^R(h_{t+1}) > K^M(h_{t+1})$.

We can now use the results of the section to gain some intuition about the short-run dynamics of innovation. Suppose radical innovation has become temporarily suboptimal, that is, $0 < K^R < K^M$. Agents will then start experimenting with technologies within the frontier. Proposition 2 shows that exploitation is never optimal, because the value of marginal innovation is strictly positive at any history. The wave of marginal innovations that follows reduces the length of the available units, as the frontier stays the same during marginal innovation. The reduction in length might then reduce the value of marginal innovation to the point of triggering a new round of radical innovation.[18] From Lemma 2, $K^R(h_t) > K^M(h_t)$ at some history h_t implies that the time-t decision maker prefers radical over marginal innovation, and the optimal size of the radical innovation $y^R(h_t) = x^R(h_t) - x_t^f$ is characterized in Proposition 3. Then, the realized outcome $f(x^R(h_t))$ determines whether radical innovation is chosen at time $t+1$ as well. For instance, suppose that the realized outcome is very high, then we know from Proposition 5 that radical innovation is again optimal provided that $f(x^R(h_t))$ exceeds the cutoff $\hat{f}(h_t)$.

In line with previous literature on neo-Schumpeterian growth, our economy witnesses growth enhancing cycles of marginal innovation followed by periods of radical innovation.[19] On the

[18] It can be shown that $\frac{dn}{dl} > 0$ from (18) in the appendix.

[19] See Jovanovic and Rob (1990) and Matsuyama (1999).

one hand, radical innovation generates a positive externality for all future generations by expanding the feasible set. The creation of an additional unit endows the society with new opportunities for costless marginal innovation, which were not available before. On the other hand, marginal innovation refines the beliefs about feasible technologies leaving new generations with a better understanding of the available opportunities. Thus, the choice of a technology creates an informational spillover which, however, is not internalized by the current generation.

5 Long-Run Dynamics of Innovation

A radical innovation has a relevant social value because it creates a positive intergenerational externality. The stock of knowledge of society increases, and new technologies become available at no cost. It is thus natural to ask whether agents will face sufficient incentives to perform radical innovation in the long-run.

Lemma 4 shows that the size of the gap might deter a decision maker from engaging in radical innovation.

LEMMA 4 *There exists* $\tilde{\rho} > 0$ *such that* $\rho_t > \tilde{\rho}$ *implies* $K^R(h_t) < 0$*, for any history* h_t*.*

The intuition follows from Proposition 4. An increase in the gap reduces the incentives to perform radical innovation through the increased appeal of simple exploitation. However, we cannot conclude directly from Lemma 4 that radical innovation is doomed in the long-run. The reason is that the choice of technologies is endogenous in our model, and thus the value of radical innovation may stay positive for a set of histories of positive measure.

We are now ready to state our main result about radical innovation.

PROPOSITION 6 *The value of radical innovation becomes negative in finite time with probability one.*

Proposition 6 provides an impossibility result on the long-run sustainability of radical innovation. Here is some intuition underlying the result. Suppose the value of radical innovation is always positive along the equilibrium path of play, contrary to Proposition 6. It must then be the case that, for any time t, there is going to be a time-t' decision maker, with $t' \geq t$, willing to perform radical innovation. This intermediate step follows from the observation that whenever radical innovation is temporarily suboptimal agents are forced to experiment inside the frontier. Since we are assuming that the value of radical innovation is positive, it follows that the sequence of marginal innovations cannot lead to the discovery of technologies that push the gap above the critical level $\tilde{\rho}$, otherwise the value of radical innovation would be negative by Lemma 4. Thus, the value of marginal innovation must necessarily decrease below the value of radical innovation after a finite number of periods. As a result, the frontier keeps expanding. As the frontier is moved further to the right, it will eventually hit a region where the gap exceeds $\tilde{\rho}$, thus reaching a contradiction.

As the frontier stops expanding, the economy witnesses a perpetual wave of marginal innovations. However, the opportunities for valuable innovation become scarce as time goes by, which leads to the following result.

Proposition 7 *The value of marginal innovation converges to zero with probability one.*

Proposition 7 says that the economy moves toward stagnation in the long-run, which was to be expected after radical innovation has stopped. Nevertheless, the value of marginal innovation is zero only in the limit, as agents keep innovating (marginally) on the equilibrium path of play. Appendix B proves the following upper bound on the value of marginal innovation, after any history h_t,

$$K^M(h_t) \leq \frac{\delta\sigma}{2(1+\delta)\sqrt{2\pi}}\sqrt{\Delta_t}$$

where Δ_t is the length of the largest finite unit with positive value of innovation at time t. The proof is completed by noticing that, as agents keep innovating marginally, Δ_t decreases over time and eventually converges to zero.

Proposition 7 only characterizes the evolution of the value of marginal innovation in the long-run, but it is silent about the limit behavior. The next result fills the gap by showing that the economy witnesses an *informational cascade*, which is defined as the convergence of the sequence of optimal technology choices.[20]

[20]The definition corresponds to the one provided in Lee (1993).

Proposition 8 *There exists a (essentially) unique technology x^* in the feasible set such that the sequence of optimal technologies converges almost surely to x^*.*

Proposition 8 exploits the equivalent interpretation of our framework as a learning problem, in which decision makers are trying to learn the underlying outcome function $f(\cdot)$. As agents try new technologies, the belief about the outcome function is updated according to Bayes' rule. Since it is well-known that the sequence of posterior beliefs follows a martingale under Bayesian updating, we can show that equilibrium beliefs converge to a limit belief $\bar{\mu}$. The next step links limit beliefs with limit behavior and establishes that any limit point of a sequence of equilibrium technologies must be optimal under the limit belief $\bar{\mu}$. Finally, suppose by way of contradiction that there exists two limit technologies x^* and x^{**}. Since they must both be accumulation points, the Brownian structure of the problem implies that they have to have different outcomes with probability one, which contradicts the fact that both technologies are optimal given $\bar{\mu}$. The details of the proof are provided in Appendix B.

6 Cost Externalities

The analysis so far has assumed that new discoveries have no effect on the cost of innovation. It seems natural to wonder whether introducing such a *cost externality* would help sustaining innovation, and thus escape the long-run stagnation identified in Proposition 6. To this end, we assume that radical innovation incurs a cost $c(x - x^f, \alpha z)$, which is now dependent on the best technology. The parameter $\alpha \geq 0$ measures the transfer of knowledge to radical innovation. In particular, the case $\alpha = 0$ embeds our previous analysis. We make the following additional assumption.

Assumption 2 *$c(\cdot, \cdot)$ is twice continuously differentiable, submodular,*[21] *and decreasing in its second component.*

Assumption 2 requires that an increase in z leads to a reduction of the cost of radical innovation. The requirement that the cost function be submodular implies that the marginal

[21] Given smoothness of c, the assumption is equivalent to c having everywhere a nonnegative cross partial derivative.

cost of radical innovation decreases as well, when the current best technology increases. Thus, a young decision maker has greater incentives to perform radical innovation when the outcome of past innovations has been successful. Since $\{z_t\}$ forms a nondecreasing sequence, the externality is persistent across all future generations.

Proposition 9 (Comparative statics) *Suppose that $\alpha > 0$. The optimal size of the radical innovation, $y^R(0, \alpha z)$, is increasing in z and in α.*

When the gap is zero, an increase in the best available technology stimulates radical innovation by decreasing the marginal cost of innovation. When the gap is positive, it is not clear anymore whether an increase in z necessarily induces more innovation because the marginal benefit and the marginal cost move in opposite directions. However, an increase in α (weakly) increases the optimal size of the radical innovation, even with a positive gap.

The following example illustrates how the short-run dynamics might change in the presence of cost externalities compared to the benchmark setting.

Example 1

Suppose the cost function has the form $c(y, \alpha z) = \int_0^y \frac{\delta\sigma}{2\sqrt{2\pi}} e^{-\frac{\alpha z^2}{2\sigma^2 s}} ds$, which satisfies Assumptions 1 and 2. To simplify the analysis, suppose also that $\alpha > 1$. Then, the first-order condition for radical innovation reduces to

$$\frac{1}{\sqrt{y}} = e^{\frac{1}{2\sigma^2 y}[(z-f(x^f))^2-\alpha z^2]} = e^{\frac{1}{2\sigma^2 y}[(1-\alpha)z^2-2f(x^f)z+f(x^f)^2]} \tag{12}$$

If $f = f(x^f) = z$, the right-hand side is increasing in y while the left-hand side is decreasing. Therefore, there is a unique nondegenerate solution, which implies a strictly positive value of radical innovation.

If instead $f < z$, the determinant of the second-order equation in square brackets is $4\alpha f^2 \geq 0$, and the second-order equation admits two solutions $z_1 = f\frac{1+\sqrt{\alpha}}{1-\alpha}$ and $z_2 = f\frac{1-\sqrt{\alpha}}{1-\alpha}$. Suppose first that $f > 0$, then $z_1 < 0 < z_2$. The first-order condition (12) always admits a positive solution for any $z > f$. When $z \in (0, z_2]$, both sides of equation (12) are decreasing in y but the right-hand side converges to $+\infty$ faster than the left-hand side, as $y \to 0$. The most interesting case occurs when $z > z_2$. In this case, the right-hand side of (12) is increasing in y. Thus, the optimal size of the radical innovation is always strictly positive when $z > \max\{z_2, f\}$. In both cases, it can be shown that an increase in z increases the size of the optimal radical innovation.

When $f \leq 0$, then $z_2 \leq 0 \leq z_1$ and a similar analysis applies. Consequently, the optimal size of the radical innovation is again strictly positive. $\square$

Example 1 shows that Proposition 6 may fail in the presence of a permanent cost externality. The intuition is that 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 a young decision maker would always set the size of radical innovation to zero.

The next result shows that stagnation can still occur, provided that the marginal cost of radical innovation doesn't approach zero as the best available technology improves.

PROPOSITION 10 *Suppose that* $\alpha > 0$ *and that* $\inf_{z>0} \frac{\partial^2 c(0,\alpha z)}{\partial y^2} > 0$*. Then, radical innovation ends in finite time with probability one.*

The additional requirement on the cost function rules out cost functions like those used in Example 1. The condition seems a mild one, because technological advancement still poses challenges at the frontier. Even if the long-run dynamics is the same with and without intergenerational cost externalities, the short-run pattern of innovation might be significantly different in the two scenarios.

7 Optimistic Beliefs

The benchmark model assumes that agents are completely ignorant about the expected payoff of technologies to the right of the frontier. What if agents are optimistic about the underlying generating process? Is a positive drift enough to sustain radical innovation?

A positive drift complicates the analysis because it creates incentives for an old decision maker to innovate as well. As before, an old decision maker prefers to exploit rather than innovate marginally. As far as radical innovation is concerned, an old decision maker faces the following problem

$$V^{O,R}(h_t) = \max_{x \in [x_t^f, +\infty)} E_{h_t}\left[f(x) - c(x - x_t^f)\right] = f(x_t^f) + \mu(x - x_t^f) - c(x - x_t^f) \tag{13}$$

In order to make the problem interesting, we need to assume that $\lim_{y \to +\infty} \frac{dc(y)}{dy} > \mu$. Otherwise, an old decision maker would push the size of radical innovation all the way to infinity. The following simple first-order condition completely characterizes optimal behavior,

$$\mu = \frac{dc(x - x_t^f)}{dy} \tag{14}$$

Let $y^{O,R}$ denote the solution to (13). A few remarks. First, the size of the optimal radical innovation of an old decision maker is independent of the history. Second, $y^{O,R}$ is positive if and only if $\frac{dc(0)}{dy} < \mu$. When $\frac{dc(0)}{dy} \geq \mu$, an old decision maker prefers to exploit, and the previous analysis repeats unaltered. Third, an old decision maker follows a simple cutoff rule, as shown in the next result.

Lemma 5 *Suppose* $\frac{dc(0)}{dy} < \mu$. *Given any history* h_t, *an old decision maker chooses*

$$x^{O,R}(h_t) = \begin{cases} y^{O,R} + x_t^f & \text{if } \rho_t \leq \xi \\ \hat{x}_t & \text{otherwise} \end{cases}$$

where $\xi = \mu y^{O,R} - c(y^{O,R}) > 0$.

The old decision maker uses the size of the gap to decide whether to innovate or not. For the decision maker, the gap measures the relative attractiveness of exploitation over radical innovation. Thus, a low gap gives high incentives to innovate.

The presence of another active agent clearly changes the short-run dynamics of the model. For instance, a young decision maker will now take into consideration that the old generation might be innovating radically before choosing his optimal technology choice. Nevertheless, radical innovation is still going to end with probability one under a mild condition on the cost function.

Proposition 11 *Suppose that* $\lim_{y \to +\infty} \frac{dc(y)}{dy} > \mu(1 + \delta)$. *Then, the value of radical innovation becomes negative in finite time with probability one, even with* $\mu > 0$ *but finite.*

The condition expressed in Proposition 11 is essentially a requirement on the third derivative of the cost function, which is implied for example by the convexity of the marginal cost. To understand the driving force underlying Proposition 11, we consider the case $\mu > \frac{dc(0)}{dy}$. When choosing the size of his innovation, a young decision maker has in mind the effect of his action today on his incentives tomorrow. This effect can be quantified in a perceived reduction of the gap from ρ_t to $\rho_t - \xi$, which results in higher incentives to perform radical innovation today. As the gap increases, though, ξ becomes negligible and eventually the relative benefit of radical innovation over exploitation falls short of the explicit cost of innovation. Thus, the young decision maker will eventually opt for marginal innovation, which is still strictly better than exploitation.

8 Conclusion

Our analysis suggests that radical innovation cannot be sustained in the long-run, when innovation relies on the short-term profitability considerations of each generation. Thus, external intervention is needed. For instance, a benevolent social planner, who internalizes the intergenerational benefits of radical innovation might find optimal to subsidize radical innovation, once the stagnation phase has been reached. Examples of public intervention in funding research are plentiful. For instance, the National Science Foundation's budget request for FY 2010 was roughly \$7 billion aimed at funding projects in areas as diverse as mathematics, engineering, and social sciences.[22] Similarly, the National Institutes of Health supervises federal funding over medical and health-related research with an impressive budget of over \$30 billion for just the fiscal year 2010.[23] These numbers seem to indicate that subsides are a relevant component of the incentives to perform innovative research. While our model corroborates this intuition, a complete analysis of the effect of subsidies on the long-run dynamics of innovation is left for future research.

[22]Source: http://www.nsf.gov/about/budget/fy2010.

[23]Source: http://www.officeofbudget.od.nih.gov.

Appendix

A Omitted Proofs

Proof of Lemma 1 We first consider a unit of finite length. Notice that

$$\max_{x\in[0,x^f]} E_{h_t}\left[f(x)+\delta\max\{f(x),z\}\right] = \max_{[x_l,x_r]\in\mathcal{P}(h_t)}\left\{\max_{x\in[x_l,x_r]} E_{[x_l,x_r]}\left[f(x)+\delta\max\{f(x),z\}\right]\right\}$$

where $\mathcal{P}(h_t)$ is the collection of finite units induced by h_t on $[0,x^f]$. We fix a finite unit u and suppose, without loss of generality, that $f(x_r)>f(x_l)$. Then, for any $x\in[x_l,x_r]$,

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

Let $g(x)=\frac{x-x_l}{x_r-x_l}$, then we can rewrite

$$f(x)-m\sim N\left(d\,g(x),g(x)(1-g(x))l\sigma^2\right)$$

If we define $k(x)=f(x)-m$ and $z'=z-m$, we get

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

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

$$\begin{aligned}E_{[x_l,x_r]}\left[f(x)+\delta\max\{f(x),z\}\right] &= (1+\delta)m+dg(x)+\delta E_{[x_l,x_r]}\max\{k(x),z'\}\\ &= (1+\delta)m+d\left\{g(x)(1+\delta-\delta\Phi)+\delta\tfrac{z'}{d}\Phi\right.\\ &\qquad\left.+\delta\sigma\sqrt{g(x)(1-g(x))}\tfrac{\sqrt{l}}{d}\phi\right\}\\ &= (1+\delta)\left\{m+d\,\bar{\eta}\left(g(x),\tfrac{\sqrt{l}}{d},\tfrac{z'}{d}\right)\right\}\end{aligned}\tag{15}$$

where

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

Finally, taking the maximum over $x \in [x_l, x_r]$, we obtain

$$\max_{x \in [x_l, x_r]} E\left[f(x) + \delta \max\{f(x), z\}\right] = (1+\delta)\left[m + d \max_{g(x)\in[0,1]} \bar{\eta}\left(g(x), \frac{\sqrt{l}}{d}, \frac{z'}{d}\right)\right]$$
$$= (1+\delta)\left[m + d\,\eta\left(\frac{\sqrt{l}}{d}, \frac{z'}{d}\right)\right]$$

If we set

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

the first part of the proposition follows immediately.

Next, suppose $l = +\infty$. The expected utility of any action $x > x^f$ is

$$E_{h_t}[f(x) - c(x - x_r) + \delta \max\{f(x), z\}] = (1+\delta)f(x^f) - c(x - x^f) + \delta E_{h_t} \max\{f(x) - f(x^f), z - f(x^f)\}$$
$$= (1+\delta)f(x^f) - c(x - x^f) + \delta E_{h_t} \max\{k(x), \rho\}$$

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

$$E_{h_t} \max\{k(x), \rho\} = E[k(x)|k(x) > \rho]\text{Prob}(k(x) > \rho) + \rho\text{Prob}(k(x) < \rho)$$
$$= \sigma\sqrt{x - x^f}\,\phi\left(\frac{\rho}{\sigma\sqrt{x - x^f}}\right) + \rho\Phi\left(\frac{\rho}{\sigma\sqrt{x - x^f}}\right)$$

Taking the supremum over $x \geq x^f$, we obtain

$$\sup_{x \geq x^f}\left\{(1+\delta)f(x^f) - c(x - x^f) + \delta\sigma\sqrt{x - x^f}\,\phi\left(\frac{\rho}{\sigma\sqrt{x - x^f}}\right) + \delta\rho\Phi\left(\frac{\rho}{\sigma\sqrt{x - x^f}}\right)\right\}$$
$$= (1+\delta)\left[f(x^f) + \sup_{y\geq 0}\frac{1}{1+\delta}\left\{-c(y) + \delta\sigma\sqrt{y}\phi\left(\frac{\rho}{\sigma\sqrt{y}}\right) + \delta\rho\Phi\left(\frac{\rho}{\sigma\sqrt{y}}\right)\right\}\right]$$

We can thus define

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

To complete the proof, let $\gamma(u,z) = f(x^f) + \psi(\rho)$.

Proof of Lemma 2 The first part of the result follows directly from the index representation provided in Lemma 1. Next, fix a finite unit u with $M \neq z$, then

$$\frac{d\eta}{dz} = \frac{\delta}{(1+\delta)d}\Phi\left(\frac{z - m - dx^*}{\sigma\sqrt{x^*(1-x^*)l}}\right)$$

where x^* is an optimal technology within u. Thus, $\frac{\partial K(u,z)}{\partial z} = \frac{\delta}{1+\delta}\Phi\left(\frac{z-m-dx^*}{\sigma\sqrt{x^*(1-x^*)l}}\right) - 1 < 0$. If, instead, u is such that $M = z$, then $\frac{\partial K(u,z)}{\partial z} = \eta\left(\frac{\sqrt{l}}{d}, 1\right) - 1$, which can be shown to be strictly positive from the argument used in Proposition 2. Then, $\frac{\partial K^M(h_t)}{\partial z} = \max_{u\in\mathcal{P}(h_t)} \frac{\partial K(u,z)}{\partial z} > 0$.

Next, we fix f and consider the effect of an increase in $z > f$ on $K^R(h_t)$. It is immediate to check that

$$\frac{\partial K^R(h_t)}{\partial z} = \frac{d\psi(z-f)}{d\rho} - 1 = \frac{\delta}{1+\delta}\Phi\left(\frac{z-f}{\sigma\sqrt{y^*}}\right) - 1 < 0$$

If, instead, $\rho = 0$, then $K^R(h_t) = \psi(0)$, and thus $\frac{\partial K^R(h_t)}{\partial z} = 0$.

Finally, if $K^R(h_t) < 0$, the above observations together with the monotonicity of $\{z_t\}$ imply that radical innovation will never be undertaken for any $t' > t$.

Proof of Proposition 1 The expected utility from any action $x > x_0^f = 0$ is

$$\begin{aligned} E_{h_t}\left[f(x) - c(x) + \delta f^+(x)\right] &= -c(x) + \delta E[f(x)|f(x) > 0]\text{Prob}(f(x) > 0) \\ &= -c(x) + \delta\left[\sigma\sqrt{x}\phi(0)\right] \end{aligned}$$

where $f^+(x) = \max\{f(x), 0\}$. The first-order condition is

$$\frac{\delta\sigma}{2\sqrt{x}}\phi(0) = \frac{dc(x)}{dx} \tag{17}$$

The right-hand side of (17) is increasing in x, while the left-hand side is strictly decreasing. However, the left-hand side is also unbounded around 0 and it converges to 0, as $x \to +\infty$. Thus, the optimal technology satisfies $0 < x_0^* < +\infty$. The SOC is

$$-\frac{\delta\sigma\phi(0)}{4x^{3/2}} - \frac{d^2c(x)}{dx^2} < 0$$

Thus, the first-order condition is necessary and sufficient.

Finally, the comparative statics follows immediately from the first-order condition.

Proof of Proposition 2 Recall that

$$\eta(k_1, k_2) = \sup_{\tilde{x}\in[0,1]} \quad \bar{\eta}(\tilde{x}, k_1, k_2) = \frac{1}{1+\delta}\left\{\tilde{x}\left(1+\delta-\delta\Phi\left(\frac{k_2-\tilde{x}}{\sigma k_1\sqrt{\tilde{x}(1-\tilde{x})}}\right)\right)\right.$$
$$\left. +\delta k_2\Phi\left(\frac{k_2-\tilde{x}}{\sigma k_1\sqrt{\tilde{x}(1-\tilde{x})}}\right) + \delta\sigma\sqrt{\tilde{x}(1-\tilde{x})}k_1\phi\left(\frac{k_2-\tilde{x}}{\sigma k_1\sqrt{\tilde{x}(1-\tilde{x})}}\right)\right\} \tag{18}$$

Suppose that $k_2 = 1$. If we differentiate $\bar{\eta}(x, k_1, 1)$ with respect to x, we obtain

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

which tends to $-\infty$ as x goes to 1. Also, $\bar{\eta}(0, k_1, 1) = \frac{\delta}{1+\delta} < 1 = \bar{\eta}(1, k_1, 1)$. We have thus shown that $\tilde{x}^*(k_1, 1) \in (0, 1)$, for any $k_1 > 0$.

Next, fix any history h_t. Consider the unit u which has z_t as one of the endpoints. Then,

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

Thus, $\gamma(u, z_t) > z_t$ proving that it is indeed strictly optimal to innovate marginally.

Next, notice that

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

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

Proof of Proposition 3 When $\rho = 0$, the first-order condition is exactly as (17), and the same result holds.

Next, suppose $\rho > 0$. The expected utility of any action $x > x^f$ is

$$\begin{aligned} & f(x^f) - c(y) + \delta\left\{f(x^f)\left(1 - \Phi\left(\frac{\rho}{\sigma\sqrt{y}}\right)\right) + \sigma\sqrt{y}\phi\left(\frac{\rho}{\sigma\sqrt{y}}\right)\right. \\ & \quad \left. + (\rho + f(x^f))\,\Phi\left(\frac{\rho}{\sigma\sqrt{y}}\right)\right\} \\ & \quad = (1+\delta) f(x^f) - c(y) + \delta\left\{\rho\Phi\left(\frac{\rho}{\sigma\sqrt{y}}\right) + \sigma\sqrt{y}\phi\left(\frac{\rho}{\sigma\sqrt{y}}\right)\right\}\end{aligned}$$

where $\rho = z - f(x^f)$. The first-order condition is then given by (11).

Proof of Proposition 4 Differentiating (11) with respect to ρ, we obtain $-\frac{\delta\rho}{2\sigma y^{3/2}}\phi\left(\frac{\rho}{\sigma\sqrt{y}}\right)$, which is strictly negative. Thus, the objective function in (16) is submodular in (y, ρ), and the result follows from the Strict Monotonicity Theorem of Edlin and Shannon (1998).

Proof of Lemma 3 Let f be the outcome from radical innovation following history h_t. The new index for marginal innovation is

$$\gamma^M(h_{t+1}) = \max \left\{ \underbrace{\max_{u \in \mathcal{P}(h_t)} \gamma(u, \max\{z_t, f\}}_{m_{h_t}(\max\{z_t, f\})}, \gamma(u'(f), \max\{z_t, f\}) \right\}$$

Case 1: $f < z_t$. Clearly, $m_{h_t}(z)$ is unaffected by changes in f, while the index for $u'(f)$ changes. In particular, let f_a be the outcome associated with the left endpoint of unit u'. If $f_a < f$, then

$$\begin{aligned}
\frac{\partial \gamma(u'(f), z)}{\partial f} &= \frac{\partial}{\partial f}\left[f_a + (f - f_a)\eta\left(\frac{\sqrt{l}}{f - f_a}, \frac{z - f_a}{f - f_a}\right)\right] \\
&= \frac{\partial}{\partial (f - f_a)}\left[f_a + (f - f_a)\eta\left(\frac{\sqrt{l}}{f - f_a}, \frac{z - f_a}{f - f_a}\right)\right] \\
&= \eta\left(\frac{\sqrt{l}}{f - f_a}, \frac{z - f_a}{f - f_a}\right) - \frac{\sqrt{l}}{f - f_a}\frac{\partial \eta}{\partial k_1}\left(\frac{\sqrt{l}}{f - f_a}, \frac{z - f_a}{f - f_a}\right) - \frac{z - f_a}{f - f_a}\frac{\partial \eta}{\partial k_2}\left(\frac{\sqrt{l}}{f - f_a}, \frac{z - f_a}{f - f_a}\right) \\
&= \frac{x^*}{1+\delta}\left[1 + \delta - \delta\Phi\left(\frac{k_2 - x^*}{\sigma k_1 \sqrt{x^*(1 - x^*)}}\right)\right] \qquad (19) \\
&> 0
\end{aligned}$$

where $x^* \in \arg\max_{x \in [0,1]} \bar{\eta}\left(x, \frac{\sqrt{l}}{f - f_a}, \frac{z - f_a}{f - f_a}\right)$, for the unit $u'(f)$. If $f_a > f$, then

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

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

$$
\begin{aligned}
(1-\tilde{x})\left[1-\frac{\delta}{1+\delta}\Phi\left(\frac{k_2-\tilde{x}}{\sigma k_1\sqrt{\tilde{x}(1-\tilde{x})}}\right)\right] &\leq \frac{1}{2}\left[1-\frac{\delta}{1+\delta}\frac{1}{2}\right] \\
&= \frac{2+\delta}{4(1+\delta)} \\
&< \frac{1}{1+\delta} \qquad (21)
\end{aligned}
$$

for any $\delta \in [0,1]$.

Case 2: $f > z_t$. We first consider a unit $u \in \mathcal{P}(h_t)$. Then,

$$
\begin{aligned}
\frac{\partial\gamma(u,f)}{\partial f} &= \frac{\partial}{\partial f}\left[m + d\eta\left(\frac{\sqrt{l}}{d}, \frac{f-m}{d}\right)\right] \qquad (22) \\
&= \frac{\partial \eta}{\partial k_2} \\
&= \frac{\delta}{1+\delta}\Phi\left(\frac{k_2 - x^{**}}{\sigma k_1\sqrt{x^{**}(1-x^{**})}}\right) \\
&\in \left(0, \frac{\delta}{1+\delta}\right)
\end{aligned}
$$

where $x^{**} \in \arg\max_{x\in[0,1]} \bar{\eta}\left(x, \frac{\sqrt{l}}{d}, \frac{f-m}{d}\right)$, for the unit u. Also, the last inequality follows from $x^{**} \in \left(\frac{1}{2}, 1\right)$, by Proposition 2.

Next, since $f > z \geq f_a$, we have

$$
\begin{aligned}
\frac{\partial\gamma(u'(f),f)}{\partial f} &= \frac{\partial}{\partial f}\left[f_a + (f-f_a)\eta\left(\frac{\sqrt{l}}{f-f_a}, 1\right)\right] \qquad (23) \\
&= \eta\left(\frac{\sqrt{l}}{f-f_a}, 1\right) - \frac{\sqrt{l}}{f-f_a}\frac{\partial\eta}{\partial k_1} \\
&= \frac{1}{1+\delta}\left[\hat{x}(1+\delta) + \delta(1-\hat{x})\Phi\left(\frac{1}{\sigma k_1}\sqrt{\frac{1-\hat{x}}{\hat{x}}}\right)\right] \\
&\in \left(\frac{\delta}{1+\delta}, 1\right)
\end{aligned}
$$

where the last inequality follows from $\hat{x} \in \left(\frac{1}{2}, 1\right)$, by Proposition 2. Comparing (22) and (23) and recalling $\delta < 1$, it is immediate to notice that the slope of the index of the new unit is always higher than the slope of the index for any old unit, whenever $f > z$.

Combining Case 1 and Case 2 shows that the two indexes intersect exactly once. The intersection occurs at an outcome $\bar{f} \le z$ if $m_{h_t}(z) \ge \gamma(u'(z), z)$, otherwise $\bar{f} > z$.

Proof of Proposition 5 The index for radical innovation at time $t+1$ is given by $\gamma^R_{t+1}(f) = f + \psi(\rho_{t+1}) = f + \psi(\max\{(z_t - f), 0\})$. So, for any $f \ge z_t$, $\gamma^R_{t+1}(f) = f + \psi(0)$, and then $\frac{\partial \gamma^R_{t+1}}{\partial f} = 1$.

Next, suppose $f < z_t$. Lemma 4 shows that there exists $\tilde{\rho}$ such that $\rho_{t+1} > \tilde{\rho}$ implies a negative value of radical innovation, that is, the optimal radical innovation is given by $x^R(h_{t+1}) = f$. Let $\tilde{f}$ be such that $z_t - \tilde{f} = \tilde{\rho}$. Then, for any $f \le \tilde{f}$, it follows that $\gamma^R_{t+1}(f) = \frac{f + \delta z_t}{1+\delta}$, which implies $\frac{\partial \gamma^R_{t+1}}{\partial f} = \frac{1}{1+\delta}$.

For $f \in (\tilde{f}, z_t)$, $\frac{\partial \gamma^R_{t+1}}{\partial f} = 1 - \frac{\delta}{1+\delta}\Phi\left(\frac{z_t - f}{\sigma\sqrt{y^*}}\right) \in \left(\frac{1}{1+\delta}, 1\right)$, where y^* is the optimal size of the radical innovation following history h_{t+1}.

Case 1: $\rho_t = 0$. Let $\bar{f}$ be the cutoff derived in Lemma 3. If $\bar{f} \ge z_t$, the index of marginal innovation is flat for any $f < z_t$. While

$$\frac{\partial \gamma^M_{t+1}}{\partial f} \in \begin{cases} \left(0, \frac{\delta}{1+\delta}\right) & \text{if } z_t \le f < \bar{f} \\ \left(\frac{\delta}{1+\delta}, 1\right) & \text{if } f \ge \bar{f} \end{cases}$$

which follows from the proof of Lemma 3. Thus, γ^R_{t+1} and γ^M_{t+1} as functions of f cross exactly once. Let $\hat{f}$ denote such intersection.

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

$$\frac{\partial \gamma^M_{t+1}}{\partial f} \in \begin{cases} 0 & \text{if } f < \bar{f} \\ \left(0, \frac{1}{1+\delta}\right) & \text{if } \bar{f} \le f < z_t \\ \left(\frac{\delta}{1+\delta}, 1\right) & \text{if } f \ge z_t \end{cases}$$

and, once again, there is a unique intersection.

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

If $\bar{f} < z_t$, the slope of the index of marginal innovation over the range $[\bar{f}, z_t)$ is given by (19), which cannot be compared with the slope of γ^R_{t+1} in an unambiguous way. Thus, there could be multiple intersections between the two indexes.

Proof of Lemma 4 The right-hand side of (11) is strictly increasing in y, under Assumption 1. The left-hand side of (11) flattens towards 0 as ρ increases, and for fixed ρ, the left-hand side converges to 0 as either $y \to 0$ or $y \to +\infty$. Also, the left-hand side is bounded above by $\frac{\delta\sigma^2\phi(1)}{2\rho}$, which converges to 0, as ρ increases. Thus, there exists $\rho' > 0$ such that $\rho > \rho'$ implies $y^R(\rho) = 0$. Next, define $\tilde{\rho} = \inf\{\rho > 0 : \rho'' > \rho \Rightarrow y^R(\rho'') = 0\}$. Clearly, $\rho > \tilde{\rho}$ implies $y^R(\rho) = 0$.

Proof of Proposition 6 We already know from Lemma 4 that there exists $\tilde{\rho} > 0$ such that $x^R(\rho) = x^f$, for any $\rho > \tilde{\rho}$. To prove the result, we need to show that

$$\text{Prob}\left(z_t - f(x_t^f) > \tilde{\rho},\ \text{for some } t \geq 0\right) = 1.$$

Without loss of generality, it suffices to focus on $\text{Prob}\left(f(x_t^R) < -\tilde{\rho},\ \text{for some } t \geq 0\right)$. Suppose by way of contradiction that radical innovation happens infinitely often, fix a sample path for which this happens, and let $\{\phi(t)\}$ be the subsequence for which $y^R_{\phi(t)} = x^R_{\phi(t)} - x^f_{\phi(t)} > 0$, for any $t \geq 0$.

Lemma 6 $x^f_{\phi(t)} \to +\infty$.

Proof. [Proof of Lemma] Suppose by way of contradiction that $x^f_{\phi(t)} \to \tilde{x} < +\infty$, then this implies that $y^R_{\phi(t)} \to 0$. By Proposition 4, there exists a further subsequence $\{\psi(t)\}$ such that the sequence $\{\rho_{\psi(t)} = z_{\psi(t)} - f(x^f_{\psi(t)})\}$ is increasing, and bounded above by $\tilde{\rho}$ from the definition of $\{\psi(t)\}$. Therefore, $z_{\psi(t)} - f(x^f_{\psi(t)}) \to \tilde{g}$. If $\tilde{g} < \tilde{\rho}$, then it follows that $y^R_{\psi(t)} \to \bar{y} > 0$, which is a contradiction. Therefore, we must necessarily have $\tilde{g} = \tilde{\rho}$, which also implies that $z_{\psi(t)} \to \tilde{z}$. This means that for sufficiently high t, the expected payoff from radical innovation is approximately equal to $f(\tilde{x}) + \delta\tilde{z}$,[24] which is strictly less than the payoff from exploitation $(1+\delta)\tilde{z}$, as $\tilde{\rho} > 0$. Thus, a young decision maker would prefer to exploit, but this contradicts the definition of the subsequence $\{\phi(t)\}$.

By Proposition 4, $\Lambda = y^R(0) \geq y^R(\rho)$, for any $\rho > 0$, and then $|x^f_{\phi(t+1)} - x^f_{\phi(t)}| \leq 2\Lambda$. Next, define

$$A_\gamma(-\tilde{\rho}) = \sup\left\{x' - x : \max_{x \leq x'' \leq x'} f(x) < -\tilde{\rho},\ \text{and } x < x' < \gamma\right\}$$

[24] Recall that the maximization problem of a young decision maker is now given by (4).

Lemma 7 *With probability one, $A_\gamma(-\tilde{\rho}) > 2\Lambda$ as $\gamma \to +\infty$.*

Proof. [Proof of Lemma] By recurrence of Brownian motion, there exists (a.s.) $\tilde{x}$ such that $f(\tilde{x}) < -\tilde{\rho}$. Therefore, there also exists $\gamma > 0$ such that $A_\gamma(-\tilde{\rho}) > 0$ a.s. Now, the result follows from the scaling property of Brownian motion.

From Lemma 6 and 7, the optimal sequence of actions will almost surely hit a region where the gap $z - f(x^f)$ exceeds $\tilde{\rho}$ and radical innovation stops, thus reaching a contradiction.

Proof of Proposition 10 Since $\inf_{z>0} \frac{\partial^2 c(0,\alpha z)}{\partial y^2} > 0$, the lower envelope generated by the family of cost functions parametrized by z is increasing. This implies that there exists

$$\tilde{\rho} = \inf \left\{ \rho \geq 0 : \ \rho' > \rho \Longrightarrow y^R(\alpha z, \rho') = 0, \ \forall \ z \geq 0 \right\}$$

such that $\rho > \tilde{\rho}$ leads to $y^R(\alpha z, \rho) = 0$, for any $z \geq 0$. Also, there exists $0 < \bar{M} < +\infty$ such that $\lim_{z \to +\infty} y^R(\alpha z, 0) < \bar{M}$, and thus we can repeat the same proof as in Proposition 6.

Proof of Proposition 11

Suppose first that $\mu < \frac{\partial c(0)}{\partial y}$. We prove that there exists $\tilde{\rho} > 0$ such that $\rho > \tilde{\rho}$ implies $y^R(\rho) = 0$. The expected utility of an old agent from radical innovation is

$$\max_{x \in [x_t^f, +\infty)} E_{h_t} \left[f(x) - c(x - x_t^f) \right] = f(x_t^f) + \mu(x - x_t^f) - c(x - x_t^f) \tag{24}$$

An old agent never pursues radical innovation in this case, while the expected utility of any technology $x > x^f$ for a young decision maker is

$$\begin{aligned} & f(x^f) + \mu y - c(y) + \delta \left\{ (f(x^f) + \mu y) \left(1 - \Phi\left(\frac{\rho - \mu y}{\sigma \sqrt{y}} \right) \right) + \sigma \sqrt{y} \phi \left(\frac{\rho - \mu y}{\sigma \sqrt{y}} \right) \right. \\ & \quad \left. + (\rho + f(x^f)) \, \Phi\left(\frac{\rho - \mu y}{\sigma \sqrt{y}} \right) \right\} \\ & \quad = (1 + \delta)[f(x^f) + \mu y] - c(y) + \delta \left\{ (\rho - \mu y) \Phi\left(\frac{\rho - \mu y}{\sigma \sqrt{y}} \right) + \sigma \sqrt{y} \phi\left(\frac{\rho - \mu y}{\sigma \sqrt{y}} \right) \right\} \end{aligned} \tag{25}$$

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

$$\mu \left(1 + \delta - \delta \Phi\left(\frac{\rho - \mu y}{\sigma \sqrt{y}} \right) \right) + \frac{\delta \sigma}{2\sqrt{y}} \phi\left(\frac{\rho - \mu y}{\sigma \sqrt{y}} \right) = \frac{dc(y)}{dy} \tag{26}$$

The left-hand side of (26) flattens towards μ as ρ increases, and for fixed ρ, the left-hand side converges to $\mu(1+\delta)$ as $y \to +\infty$, and to μ as $y \to 0$. Also, the left-hand side is bounded above by

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

which converges to $\mu(1+\delta)$, as ρ increases. Thus, $\lim_{y\to+\infty} \frac{dc(y)}{dy} > \mu(1+\delta)$ implies that there exists $\tilde{\rho} > 0$ such that $\rho > \tilde{\rho}$ implies $y^R(\rho) = 0$.

Next, if we differentiate (26) with respect to ρ, we obtain

$$-\frac{\delta\phi}{2\sigma\sqrt{y}}\left[\mu + \frac{\rho}{y}\right] < 0, \quad \forall\, \rho, y > 0$$

By Monotone Comparative statics, we can conclude that $y^R(0) \geq y^R(\rho)$, $\forall\, \rho > 0$. We can thus repeat the same argument of Proposition 6 to conclude that radical innovation ends in finite time a.s.

Next, suppose $\mu > \frac{dc(0)}{dy}$. Lemma 5 implies that $y^{O,R} > 0$ if and only if $\rho < \xi$. Thus, we assume without loss of generality that $\rho_t > \xi$, so that an old agent doesn't innovate today and then the expected utility today of any $x > x_t^f$ for a young decision maker is simply

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

The first-order condition is

$$\mu\left[1 + \delta - \delta\Phi\left(\frac{\rho_t - \xi - \mu y}{\sigma\sqrt{y}}\right)\right] + \frac{\delta\sigma}{2\sqrt{y}}\phi\left(\frac{\rho_t - \xi - \mu y}{\sigma\sqrt{y}}\right) = \frac{\partial c(y)}{\partial y} \tag{27}$$

Notice that the right-hand side is always at least μ. Since $\rho_t - \xi > 0$, then the right-hand side converges to μ as $y \to 0$, and to $\mu(1+\delta)$ as $y \to +\infty$. Therefore, there always exists a nondegenerate and finite solution to the first-order equation (27). Also, under $\lim_{y\to+\infty} \frac{dc(y)}{dy} > \mu(1+\delta)$, the solution is unique for high values of ρ. In this case, as ρ increases, $y^R(\rho)$ approaches y^{OR}, that is, the optimal size of the innovation for a young decision maker converges to the optimal size for an old decision maker. This follows from the fact that the first-order condition is approximately equal to $\mu \approx \frac{\partial c(y)}{\partial y}$. Thus, for sufficiently high ρ, the maximized expected utility is approximately equal to

$$(1+\delta)(f(x_t^f)+\xi)+\delta\left[\mu y^{O,R}\left(1-\Phi\left(\frac{\rho_t-\xi-\mu y^{O,R}}{\sigma\sqrt{y^{O,R}}}\right)\right)\right.$$
$$\left.+\sigma\sqrt{y^{O,R}}\phi\left(\frac{\rho_t-\xi-\mu y^{O,R}}{\sigma\sqrt{y^{O,R}}}\right)+(\rho_t-\xi)\Phi\left(\frac{\rho_t-\xi-\mu y^{O,R}}{\sigma\sqrt{y^{O,R}}}\right)\right]$$
$$\approx f(x_t^f)+\xi+\delta z_t$$

Since we assumed $\rho_t > \xi$, it follows that $f(x_t^f)+\xi+\delta z_t < (1+\delta)z_t$. This means that the only nondegenerate candidate for radical innovation gives an expected payoff which is lower than what the young decision maker could get by simply exploiting. Thus, there exists $\tilde{\rho} > 0$ such that the young decision maker prefers exploitation for any gap greater than $\tilde{\rho}$.

In order to apply the same steps of Proposition 6, we still need to show that there is an upper bound on the size of the innovations that agents might want to undertake. First, if an old decision maker doesn't innovate today, the value function of a young decision maker is submodular in y and ρ, which shows that $y^R(0) \geq y^R(\rho)$, for any $\rho > 0$. Next, suppose an old decision maker innovates today. This means that the outside option tomorrow is higher in expectation, which reduces the young decision maker's incentives to perform radical innovation today. Formally, if the young decision maker decides to innovate today as well, and he chooses $x_t^f < x \leq x_t^f + y^{OR}$, then his innovation clearly has y^{OR} as an upper bound. Otherwise, the expected payoff associated with any $x > x_t^f + y^{OR}$ is

$$E[f(x)-c(x-x_t^f)]$$
$$+\delta\Big\{E[f(x)+\xi|f(x)\geq\max\{z_t,f(x_t+y^{OR})\}-\xi]\ \mathrm{Prob}(f(x)\geq\max\{z_t,f(x_t^f+y^{OR})\}-\xi)$$
$$+E[\{z_t,f(x_t^f+y^{OR})\}|f(x)<\max\{z_t,f(x_t+y^{OR})\}-\xi]\ \mathrm{Prob}(f(x)<\max\{z_t,f(x_t^f+y^{OR})\}-\xi)\Big\}.$$

Submodularity implies that the optimal size of the radical innovation is largest when the gap is zero. This shows the existence of an upper bound M on the size of any radical innovation, uniform over all histories.

B Preliminary Results for Propositions 7 and 8

Since we are assuming that radical innovation has already stopped, the feasible region can be normalized to the interval $X = [0,1]$. The set of sample paths is the space of continuous functions from $(0,0)$ to $(1,d)$ with $d \geq 0$, which we denote by Θ. Since Θ is unbounded, it is convenient to apply a monotone and continuous transformation $F : \mathbb{R} \to \mathbb{R}$, with F strictly increasing, $F(0) = 0$, $F(1) = d$ and $\lim_{|y| \to \infty} |F(y)| \leq \Lambda$. We can thus define the new space $\tilde{\Theta} = F(\Theta)$, which is a subset of $C[0,1]$ as well. Without loss of generality, we can work with $\tilde{\Theta}$.

We can see our problem as a learning problem involving a stochastic process with an unknown parameter $\theta \in \tilde{\Theta}$. The initial belief μ_0 over $\tilde{\Theta}$ is represented by the law of the Brownian bridge, given $(0,0)$ and $(1,d)$. The underlying true parameter is denoted by $\theta_0 \in \tilde{\Theta}$.

For any $(x,\theta) \in X \times \tilde{\Theta}$, we can also define a conditional density function over outcomes $g(y|x,\theta) = \mathbf{1}_{\theta(x)}(y)$, for any $y \in Y = [-\Lambda, \Lambda]$, which is the Dirac measure. Let μ_t be the belief at time t and (x_t, y_t) the current action with the corresponding outcome. We can define the belief at time $t+1$ as follows

$$\mu_{t+1}(A) = \frac{\int_A g(y_t|x_t,\theta) d\mu_t}{\int_{\tilde{\Theta}} g(y_t|x_t,\theta) d\mu_t}$$

where A is a Borel subset of $\tilde{\Theta}$, and the integrals are to be interpreted in the Stieltjes sense.

The function $r : [-\Lambda, \Lambda]^2 \to \mathbb{R}$ specifies the reward following the observation of an outcome y, when the best known outcome is z. In our framework, we can restrict attention to the function $r(y,z) = y + \delta \max\{y,z\}$. For convenience, we define the payoff function

$$\begin{aligned} u(x,\mu,z) &= \int_{\tilde{\Theta}} \left[\int_{\mathbb{R}} r(y,z) g(y|x,\theta)\, dy \right] d\mu \\ &= \int_{\tilde{\Theta}} r(\theta(x), z)\, d\mu \end{aligned}$$

The following technical lemma will prove to be useful in our analysis.

Lemma 8 *$u(x,\mu,z)$ is continuous over $[0,1] \times \Delta(\tilde{\Theta}) \times [-\Lambda, \Lambda]$.*

Proof. Let $\{(x_n, \mu_n, z_n)\}$ be a sequence from $[0,1] \times \Delta(\tilde{\Theta}) \times [-\Lambda, \Lambda]$ which converges to $(x,\mu,z) \in [0,1] \times \Delta(\tilde{\Theta}) \times [-\Lambda, \Lambda]$. Then,

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

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

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

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

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

Recall that X is compact and every $\theta \in \tilde\Theta$ is continuous, therefore it is also uniformly continuous. Fix $\epsilon > 0$ and define

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

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

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

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

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

Since ϵ and η are arbitrary, this completes the proof.

We can thus rewrite the maximization problem of a young decision maker as

$$V(\mu, z) = \max_{x \in [0,1]} u(x, \mu, z), \quad \forall\, (\mu, z) \in \Delta(\tilde{\Theta}) \times [-\Lambda, \Lambda]$$

Let $X^*(\mu, z) = \text{argmax}_{x \in [0,1]}\, u(x, \mu, z)$, which is well-defined by Lemma 8, and also upper-hemicontinuous as a consequence of the Maximum Theorem. The analysis proceeds by restricting attention to a selection $x^*(\cdot)$ from the correspondence $X^*(\cdot)$.

LEMMA 9 *There exists a measurable selection $x^*(\cdot)$ from $X^*(\cdot)$.*

Proof. It is a consequence of the upper-hemicontinuity of $X^*(\cdot)$, and the Kuratowski-Ryll-Nardzewski Selection Theorem.[25]

In order to talk about convergence of beliefs, we need to be more precise about the underlying probability space.[26] Let $\Gamma = [0,1] \times [-\Lambda, \Lambda]$ and $\mathcal{H} = \prod_{t=0}^{\infty} \Gamma$. Also, define the family of t-cylinder sets $\tilde{\mathcal{H}}_t = \left\{C \subseteq \mathcal{H} :\ C = D \times \prod_{s=t+1}^{\infty} \Gamma,\ \text{for } D \text{ Borel subset of } \prod_{s=1}^{t} \Gamma\right\}$, and $\tilde{\mathcal{H}} = \bigvee_{t=0}^{\infty} \tilde{\mathcal{H}}_t$. For any $\theta \in \tilde{\Theta}$ and $C \in \tilde{\mathcal{H}}_t$, we can define the probability measure $\mathrm{P}_\theta(C) = \mathrm{P}_\theta\left(\{(x^*(\mu_s), y_s)\}_{s=0}^{t} \in D\right)$ where D is the base of the cylinder C. Also notice that P_θ is measurable in θ as a consequence of the measurability of the density $g(\cdot)$, and the optimal selection $x^*(\cdot)$. We can thus define the measurable space $(\tilde{\Theta} \times \mathcal{H}, \sigma(\mathcal{B}(\tilde{\Theta}) \times \tilde{\mathcal{H}}))$ and the measure $\tilde{P}$, which is the extension of $P(B \times C) = \int_B \mathrm{P}_\theta(C)\, d\mu_0$, for any $B \in \mathcal{B}(\tilde{\Theta})$ and $C \in \tilde{\mathcal{H}}$.

[25] See Theorem 17.13 in Aliprantis and Border (1999).

[26] The construction of the probability space follows the steps outlined in Easley and Kiefer (1988).

We can now state the following result,

Lemma 10 *Given $x^*(\cdot)$, $\{\mu_t\}$ is a martingale.*

Proof. It is well-known. For example, see Easley and Kiefer (1988).

Also, beliefs converge in the long-run.

Theorem 2 *There exists a $\tilde{\mathcal{H}}$-measurable random variable $\bar{\mu}$ such that $\mu_t \to \bar{\mu}$ $\tilde{P}$-a.s.*

Proof. It is a consequence of the Martingale Convergence Theorem. The proof is identical to Theorem 4 in Easley and Kiefer (1988).

Recall that in our framework, not every belief in $\Delta(\tilde{\Theta})$ is admissible. So, define

$$\Delta^f(\tilde{\Theta}) = \left\{\mu \in \Delta(\tilde{\Theta}) : \exists\, t \text{ and } h \in \mathcal{H} \text{ s.t. } \mu = \mu_{0|h_t}\right\}$$

where $\mu_{0|h_t}$ denotes Bayesian updating of the prior belief μ_0 following the truncated history h_t. Also, define

$$\Delta^\infty(\tilde{\Theta}) = \left\{\mu \in \Delta(\tilde{\Theta}) : \exists\, h \in \mathcal{H} \text{ s.t. } \mu_t \to \mu\ \tilde{P} - \text{a.s., where } \mu_t = \mu_{0|h_t}, \forall\, t\right\}$$

Next, take $\mu \in \Delta^f(\tilde{\Theta})$. By definition, there exists t and $h \in \mathcal{H}$ such that $\mu = \mu_{0|h_t}$. Notice that the history is not necessarily unique. However, the only thing that matters is the partition induced over $[0,1]$ together with the corresponding outcomes. Hence, the history is unique in the sense that every history that leads to the same partition leads to the same measure under Bayesian updating. Without loss of generality, we can thus focus on histories h_t such that $0 \leq x_0 \leq x_1 \leq ... \leq x_{t-1} \leq 1$, which will be referred to as *reduced histories.*

We define a function $Z : \Delta^\infty(\tilde{\Theta}) \to \mathbb{R}$ such that

1. if $\mu \in \Delta^f(\tilde{\Theta})$, then $Z_t(\mu) = f\left(\hat{x}_t(h_t)\right)$, where h_t is the unique reduced history corresponding to μ.

2. Consider now $\mu \in \Delta^\infty(\tilde{\Theta}) \backslash \Delta^f(\tilde{\Theta})$ and let $h \in \mathcal{H}$ be a history compatible with μ. We can then define the sequence $\{z_t = Z(\mu_{0|h_t})\}$, which is nondecreasing and bounded. Hence, it admits a unique limit $\hat{z}$ and we can define $Z(\mu) = \hat{z}$.

Since there might be many histories compatible with a belief $\mu \in \Delta^f(\tilde{\Theta})$, we need to show that the function $Z(\cdot)$ is well-defined.

LEMMA 11 *Fix $\mu \in \Delta^\infty(\tilde{\Theta})\backslash\Delta^f(\tilde{\Theta})$. Let $h, h' \in \mathcal{H}$ with $h \neq h'$ be two histories compatible with μ. Then, $Z(\mu|h) = Z(\mu|h')$.*

Proof. Suppose by way of contradiction that $Z(\mu|h) > Z(\mu|h')$. Then, by the properties of the Brownian Bridge, the limit beliefs associated with h and h' must be different, thus contradicting the fact that h and h' are both compatible with μ. Similarly, it cannot be $Z(\mu|h) < Z(\mu|h')$ either. Thus, we are left with $Z(\mu|h) = Z(\mu|h')$, which completes the proof.

The next steps show that convergence of beliefs implies convergence of behavior.

DEFINITION 3 *Given a sequence of actions $\{x_t\}_{t=0}^\infty$, let $\mathcal{M}(\{x_t\})$ be the set of its limit points.*

THEOREM 3 *Fix $\mu \in \Delta^\infty(\tilde{\Theta})$ and let $\{x_t = x^*(\mu_t, Z(\mu_t))\}$ be the corresponding sequence of optimal actions. Then,*

$$x \in \mathcal{M}(\{x_t\}) \Longrightarrow x \in \underset{\tilde{x}\in[0,1]}{\operatorname{argmax}}\, u(\tilde{x}, \mu, Z(\mu))$$

and $V(\mu, Z(\mu)) = (1+\delta)Z(\mu)$.

Proof. Let $x \in \mathcal{M}(\{x_t\})$. Thus, there exists a convergent subsequence $\{x_{t_k}\}$. By construction, $u(x_{t_k}, \mu_{t_k}, Z(\mu_{t_k})) \geq u(x', \mu_{t_k}, Z(\mu_{t_k}))$, for any $x' \in [0,1]$. Taking limits, we arrive at $u(x, \mu, Z(\mu)) \geq u(x', \mu, Z(\mu))$.

Next, as k increases, the length $x_r - x_l$ of the intervals containing x shrinks to zero. By the continuity of the path of a Brownian motion, $f(x_r) - f(x_l) \approx 0$, and the variance over $[x_l, x_r]$ is of the order $o(|x_r - x_l|)$. Thus, for t sufficiently large, it follows that

$$E_{[x_l^t, x_r^t]}\left[f(x) + \delta \max\{f(x), z_t\}\right] \approx (1+\delta)f(x_l^t) + \delta(z_t - f(x_l^t))$$

Since x_l^t converges to x and $f(\cdot)$ is continuous, by the limit optimality of x, we get $f(x_l^t) \to z(\mu)$. Therefore, $V(\mu, z(\mu)) = (1+\delta)z(\mu)$.

Proof of Proposition 7

LEMMA 12 *If $\gamma(u,z) < z$ for some $z > 0$, then $\gamma(u,z') < z'$ for any $z' > z$.*

Proof. The condition $\gamma(u,z) < z$ can be rewritten as

$$\eta\left(\frac{\sqrt{l}}{d}, \frac{z-m}{d}\right) < \frac{z-m}{d}$$

It is easy to show that $\frac{\partial \eta}{\partial z} = \frac{\delta}{(1+\delta)d}\Phi < \frac{1}{h}$, and the result follows.

LEMMA 13 *For any unit u of finite length, there exists $z = z(u)$ such that $\gamma(u,z) < z$.*

Proof. By (18), for $z > f(x^r)$ sufficiently high, $\eta(k_1, k_2) \approx \frac{1+\delta k_2}{1+\delta}$, which implies $\gamma(u,z) \approx \frac{f(x_r)+\delta z}{1+\delta} < z$.

Next, define $\mathcal{I}_t = \{u : K(u, z_t) \leq 0\}$ as the family of units on which innovation has ceased.

LEMMA 14 *$\bigcup_{u\in\mathcal{I}_t} u \uparrow [0,1]$, as $t \to +\infty$.*

Proof. Monotonicity follows immediately from Lemmas 12 and 13, together with the monotonicity of $\{z_t\}$. Suppose by way of contradiction that $\lim_{t\to+\infty} \bigcup_{u\in\mathcal{I}_t} u \neq [0,1]$, then there exists at least one unit $\tilde{u}$ of positive length such that $\gamma(\tilde{u}, z(\mu)) > z(\mu)$, by Proposition 2. Thus, by innovating in $\tilde{u}$, a decision maker can guarantee himself an expected payoff strictly higher than $(1+\delta)z(\mu)$, but this contradicts Theorem 3.

LEMMA 15 *For any history $h_t \in \mathcal{H}_t$,*

$$K^M(h_t) \leq \frac{\delta\sigma}{2(1+\delta)\sqrt{2\pi}}\sqrt{\Delta_t} \tag{28}$$

where $\Delta_t = \max\{l : \exists\, u \in [0,1]\backslash\mathcal{I}_t$ of length $l\}$ is the length of the largest unit with positive value of innovation at time t.

Proof. Fix a unit $u = [I, m, M]$ of finite length. It is immediate to notice that the value function for that interval is bounded above by the value function of the unit $u' = [I, z, z]$, which is given by

$$\max_{x\in[x_l,x_r]} E_{[x_l,x_r]}\left[f(x) + \delta \max\{f(x), z_t\}\right] = (1+\delta)\left[z_t + \frac{\delta\sigma\sqrt{x_r - x_l}}{2(1+\delta)\sqrt{2\pi}}\right]$$

Thus, the value of innovation of u is bounded by $\frac{\delta\sigma}{2(1+\delta)\sqrt{2\pi}}\sqrt{l}$, where l is the length of u. By definition, any $u \in \mathcal{I}_t$ has $K(u, z_t) \leq 0$, and the result follows.

Finally, Lemma 14 and Proposition 2 imply that $\Delta_t \to 0$, which completes the proof.

Proof of Proposition 8 We start with some preliminary results.

LEMMA 16 *Consider a unit u with $l > 0$, such that $\gamma(u, z) < z$, for some $z > 0$. Then, with probability 1, either $[0, x_l]$ or $[x_r, 1]$ is visited only finitely many times.*

Proof. Consider the innovation problem with initial domain $[0, x_l]$, and let $Z_l = Z(\mu)$ be the limit value of the best known outcome, which exists by Theorem 3. Similarly, we can define Z_r over $[x_r, 1]$. Z_l and Z_r are two random variable which are different a.s.

Next, suppose by way of contradiction that both $[0, x_l]$ and $[x_r, 1]$ are visited infinitely often. Given the independence of the distributions over the two subintervals, it must be the case that the subsequence visiting each subinterval coincides with the sequence of optimal actions from the corresponding subproblem, in which only that subinterval is available. Thus, the maxima on each subinterval in the more general problem must equal Z_l and Z_r. Since $Z_l \neq Z_r$ a.s., a decision maker will eventually prefer one subinterval over the other, which leads to a contradiction.

LEMMA 17 *Fix a unit u. For each $0 < q < \frac{1}{2}$, there exists a subinterval $[x_l', x_r'] \subseteq [x_l, x_r]$, with length at least $q(x_r - x_l)$, and an a.s. finite stopping time τ, such that $[x_l', x_r']$ is part of unit u' at time τ, with $\gamma(u', z_\tau) < z_\tau$.*

Proof. Suppose the lemma is false. Then, all subintervals of $[x_l, x_r]$ of length at least $q(x_r - x_l)$ are visited infinitely often. This implies that each subinterval contains a point from $\mathcal{M}(\{x_t\})$, and then we can get arbitrarily close to $z(\mu)$. However, this happens with probability zero.

LEMMA 18 *Suppose $x \in \mathcal{M}(\{x_t\})$. Then, for any $\epsilon > 0$ and $T \geq 0$, there exist $t, t' \geq T$ such that $x_t \in (x - \epsilon, x)$ and $x_{t'} \in (x, x + \epsilon)$.*

Proof. Suppose the claim is false. Without loss of generality, let $w = \min\{x_t : x_t > x,\ t \geq 0\} > x$, which is well-defined, and let u be the unit with interval $[x, w]$. In the limit, $\gamma(u, z(\mu)) >$

$z(\mu)$ by the proof of Proposition 2. Then, a decision maker would find it profitable to innovate in $[x, w]$, because the highest payoff in the limit is $(1+\delta)z(\mu)$. We have thus reached a contradiction.

Now, suppose by way of contradiction that there exists $x, y \in \mathcal{M}(\{x_t\})$, with $x < y$. We consider the interval $[x, y]$. By Lemma 18, we can always find two known actions x' and y' such that $x < x' < y' < y$. By Lemma 17, there exists a subinterval $[x'', y''] \subseteq [x', y']$, as well as an a.s. finite stopping time τ, such that $[x'', y'']$ is not visited after τ. Thus, Lemma 16 implies that either $[0, x'']$ or $[y'', 1]$ is visited only finitely many time, but this contradicts the assumption that both x and y are accumulation points.

References

Acemoglu, D., G. Gancia, and F. Zilibotti, 2010, "Competing Engines of Growth: Innovation and Standardization," *Journal of Economic Theory*, p. forthcoming.

Aghion, P., P. Bolton, C. Harris, and B. Jullien, 1991, "Optimal Learning by Experimentation," *Review of Economic Studies*, 58, 621–654.

Aliprantis, C., and K. Border, 1999, *Infinite Dimensional Analysis*. Springer, 2nd, ed.

Bala, V., and S. Goyal, 1998, "Learning from Neighbours," *Review of Economic Studies*, 65, 595–621.

Banerjee, A., 1992, "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 107, 797–817.

Bikhchandani, S., D. Hirshleifer, and I. Welch, 1992, "A Theory of Fads, Fashion, Custom, and Cultural Change as Information Cascades," *Journal of Political Economy*, 100, 992–1026.

Billingsley, P., 1968, *Convergence of Probability Measures*. John Wiley and Sons, New York, NY.

Callander, S., 2009, "Searching for Good Policies," mimeo.

Easley, D., and N. Kiefer, 1988, "Controlling a Stochastic Process with Unknown Parameters," *Econometrica*, 56, 1045–1064.

Edlin, A., and C. Shannon, 1998, "Strict Monotonicity in Comparative Statics," *Journal of Economic Theory*, 81, 201–219.

Garicano, L., and E. Rossi-Hansberg, 2009, "Organizing Growth," *Journal of Economic Theory*, p. forthcoming.

Griliches, Z., 1992, "The Search for R&D Spillovers," *Scandinavian Journal of Economics*, 94, S29–S47.

Hall, B., 1996, "The Private and Social Returns to Research and Development," *Technology, R&D, and the Economy*, Brookings Institution and the American Enterprise Institue, Washington, DC, 140–83.

Jones, C., and J. Williams, 1998, "Measuring the Social Return to R&D," *Quarterly Journal of Economics*, 113, 1119–1135.

Jovanovic, B., and Y. Nyarko, 1996, "Learning by Doing and the Choice of Technology," *Econometrica*, 64, 1299–1310.

Jovanovic, B., and R. Rob, 1990, "Long Waves and Short Waves: Growth Through Intensive and Extensive Search," *Econometrica*, 58, 1391–1409.

Kauffman, S., L. Lobo, and W. Macready, 2000, "Optimal Search on a Technology Landscape," *Journal of Economic Behavior & Organization*, 43, 141–166.

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

Matsuyama, K., 1999, "Growing Through Cycles," *Econometrica*, 67, 335–347.

McLennan, A., 1984, "Price Dispersion and Incomplete Learning in the Long Run," *Journal of Economic Dynamics and Control*, 7, 331–347.

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

Strulovici, B., 2010, "An Index Policy for Experimentation with Correlated Actions," Working Paper, Northwestern University.