Innovation and Competition on a Rugged Technological Landscape*

Steven Callander[†] Nicolas Lambert[‡] Niko Matouschek[§]

June 7, 2022

Abstract

We develop a model of spatial competition in which the quality of a product is learned only after it is introduced to the market. Firms enter sequentially, choosing whether to innovate beyond the frontier and outside the scope of the existing market, or to nestle in a niche between existing products. The uncertainty about a new product's quality depends on this choice and increases in the degree of horizontal differentiation from existing products. Innovation in this market is irregular with frequent changes of direction and cycles between frontier and niche innovation. We show how the ruggedness of the technological landscape itself deters innovation, generating less product differentiation, narrower markets, and less entry than in a world of certainty. We develop and explore numerically a targeted policy intervention that encourages innovation when it ends prematurely. The interventions are short in duration but can restart self-sustaining innovation, generating large returns in welfare.

*For helpful comments we thank Daron Acemoglu, Dana Foarta, Ben Jones, Sarit Markovich, Sara Moreira, Spencer Pantoja, Bruno Strulovici, Takuo Sugaya, Can Urgun, and Leeat Yariv.

†Stanford University, Stanford Graduate School of Business; sjc@stanford.edu.

‡Massachusetts Institute of Technology, Department of Economics; lambertn@mit.edu.

§Northwestern University, Kellogg School of Management; n-matouschek@kellogg.northwestern.edu

1 Introduction

Research on innovation has long recognized the importance of new products and the exit of old products as driving market evolution and economic growth (e.g., Romer, 1990; Wollmann, 2018). The novelty of new products varies widely, however. Some are barely tweaks on existing products, whereas others represent breakthrough innovations that depart radically from what has come before.

The breadth of novelty is clear in the pharmaceutical industry. Krieger, Li and Papanikolaou (2021) characterize at the molecular level the rich variety in new drugs, from 'me too' drugs that are close yet imperfect imitations of existing drugs, to radical innovations that target different illnesses and segments of the patient population. These differences are not happenstance, and are driven by deliberate choices of the firms.[1]

The strategic problem for a firm, therefore, is not only whether to innovate but how to innovate. Should the firm innovate incrementally or should it innovate boldly? Should it innovate within the boundaries of the existing market or outside? How will these choices affect the quality of the product produced and the nature of market competition? These questions are the essence of business strategy and have been popularized by management theorists as the choice between a red ocean strategy—of incremental change and intense competition *within* a market—versus the blue ocean strategy of pursuing new customers and open space *outside* of the market boundaries (Kim and Mauborgne, 2004).

These questions also matter for social welfare and public policy. On one hand, 'me too' drugs do improve on existing products, even if incrementally, and boost consumer welfare through heightened price competition (DiMasi and Faden, 2011). On the other hand, radically new products help patients who otherwise would not be well treated, and open up opportunities for further innovations (Krieger, Li and Papanikolaou, 2021). It is important to understand, therefore, how the different types of innovations contribute to welfare, what is the optimal balance between incremental and radical innovation, and how government policy can be used to improve on market outcomes.

In this paper, we introduce a framework that captures the strategic problem firms face and we explore its implications for innovation, competition, and industry dynamics. A key element of our approach is that we allow for a full continuum of horizontally differentiated products. This enables us to capture a sense of space and distance in the decision to innovate and, simultaneously, in the degree of competition between firms. In deciding where to locate, each firm chooses how much to differentiate from existing products and also the type of their innovation. They can innovate in a *niche*, nestling between products and competing for existing customers—a red ocean strategy—or they can innovate beyond the *frontier*, outside

[1] See Granja and Moreira (2021) for evidence from financial services.

the boundaries of existing firms, searching out new customers in the untrammeled blue ocean where risk is high but competition less intense.

Formally, our model represents a melding of innovation with Hotelling's (1929) classic model of spatial competition. The desire to find space to compete represents Hotelling's lasting insight that firms differentiate to soften competition.[2] To that, we add the observation that differentiation also means innovating and trying something new. Differentiation and innovation go hand in hand. To capture the uncertainty of trying something new, we overlay on the Hotelling line the realized path of a Brownian motion, where the outcome of the path represents a product's quality. Quality is revealed only for products that have been introduced to the market, and uncertainty about new products increases in their novelty—that is, in the distance an innovation is from existing products.[3]

For this environment we analyze a dynamic model of market competition. In each period a new firm is given the opportunity to enter the market. If it does so, it profits from its innovation in competition with incumbent firms, although this ability is short-lived as its product is quickly imitated by other firms. Over time, the new entrants fill out and expand the scope of the market, continuing until further entry is no longer profitable.

In this setting, Hotelling's intuition is thrown into stark relief. In Hotelling's world—a world without innovation—the desire to soften competition leads to maximal differentiation. Each new firm expands the frontier of the market, entering beyond the boundaries of existing products where consumers are untouched and competition is nonexistent.[4] A blue ocean strategy is dominant. In Hotelling's world, there is no benefit to competing in a niche when a firm can always escape competition by locating beyond the frontier.

The logic of entry and the evolution of market structure are different in a market with innovation. Even if uncertainty does not directly affect the preferences of the firms, it changes the landscape on which they compete. Some innovations succeed while others fail. Thus, an entering firm faces an uneven landscape, where the likelihood that a new product is of high quality varies depending on where it is located.

We show that the ruggedness of the technology landscape leads to less differentiation and, in particular, reversals in both the direction and type of innovation. The blue ocean of frontier-expanding innovation may now offer worse prospects for a high quality product. When this occurs, firms reverse the direction of innovation, turning instead to niche innovation, trading

[2]In this sense, one may interpret Hotelling as always having been about innovation, albeit without any uncertainty in product quality. We adopt the view that innovation is inseparable from uncertainty.

[3]Our definition of innovation follows Rogers (1962, p.475): "An idea, practice, or object that is perceived as new by an individual or other unit of adoptions." It does not require that the outcome be a success.

[4]We assume consumers are arrayed uniformly across the entire real line, allowing us to avoid edge effects and to study the choice between niche and frontier innovation. This formulation is due to Lancaster (1979). Salop's (1979) circle avoids edge effects, but bounds frontier-innovation as the market can be covered with a finite number of firms.

off heightened competition for the better prospects of a high quality product.

A core insight of our model is to show how this trade-off changes over time. Firms turn to niche innovation as it is more attractive than the frontier, yet the appeal of niche innovation sews the seeds of its own demise. As more firms exploit niches, those niches become crowded and competition becomes more intense. The red ocean becomes more red. This renders the frontier relatively appealing once again, and firms turn back to the blue ocean, expanding the scope of the market and creating new niches. Thus, competition shapes innovation and innovation, in turn, shapes competition. This can lead to cycles in innovation as the market develops and matures, moving with the randomness of innovation itself.

Our model also delivers a rich set of patterns on product quality and the life cycle of firms. We uncover a novel anti-differentiation force in product quality. When firms innovate in a niche, they prefer balanced competition in which their competitors are evenly matched rather than where one is strong and the other weak. We show that this preference is so strong that a firm will deliberately choose an innovation with lower expected quality and with less horizontal differentiation if it means competition is balanced. This preference generates an endogenous clustering in firm quality that is not coincidental but rather by deliberate intent.

Schumpeter's (1942) famous insight was that the replacement of the old with the new, of existing products and firms with new entrants, is the essence of innovation. Our model captures this process of creative destruction. In the model, some firms fail and disappear immediately, whereas others find a foothold in the market only to be disrupted later by a new innovation, and some firms survive through to the point when innovation ends and the market stabilizes. We explore the dynamics of disruption numerically, characterizing the life cycle of a typical firm. We show which types of innovation are more likely to disrupt the market and show how the rate of disruption varies in the complexity of the technological landscape and in the intensity of market competition.

Innovation in our model ends inefficiently early, deterred by the ruggedness of the technological landscape itself. Innovation thrives at the peaks of the landscape, but in the valleys it can get stuck. This represents a market failure as it is typically socially efficient for innovation to continue even though it is not in the interests of a single firm. We investigate a targeted policy intervention that encourages innovation when it is stuck. Although the ruggedness of the landscape is why innovation stops, it also implies that moving innovation out of a valley can have a large and lasting impact. We show how short interventions, often only for a single product, can restart innovation such that it is self-sustaining thereafter. These targeted interventions generate returns in social welfare that are many multiples of the cost.

Related Literature

Hotelling's seminal model of spatial competition has fundamentally changed our understanding of market competition. Although Hotelling's model is about trying new locations or products, it has not, to the best of our knowledge been used to model innovation *per se*.[5] Hotelling implicitly assumes identical quality across all products and focuses exclusively on horizontal differentiation. We add a rugged technological landscape that captures both horizontal and vertical product differentiation, combined with uncertainty over the landscape such that firms only learn about quality through experience.

We follow Lancaster (1979) in modeling an unbounded space of products and consumers, and Prescott and Visscher (1977) in supposing that firms locate sequentially and are fixed in their locations thereafter. This formulation resonates with the organizational sociology view that firms are inertial and that market evolution is predominantly through selection (Hannan and Freeman, 1977).[6] We differ from Prescott and Visscher (1977) in presuming firms are focused on the period of entry. This short-sightedness follows the innovation literature in supposing that above-normal returns of successful innovation are short-lived and aids considerably in the tractability of the model.

An alternative approach to horizontal differentiation is the approach of Chamberlin (1933). This tradition produced the workhorse model of Dixit and Stiglitz (1977), which has been applied to innovation in the influential growth model of Romer (1990). As impactful as this line of work has been, it obscures the micro-foundation of the innovation process itself. As Lancaster (1990, p.194) remarks, "An important limitation on the Dixit-Stiglitz and other neo-Chamberlinian models is that firms make no product choice—it is as though each firm, as it enters the group, is assigned a product by random choice (without replacement) from an urn containing blueprints for all possible products."

The innovation literature beyond Chamberlin (1933) focuses on vertical differentiation, such as in models of dynamic and strategic R&D of Reinganum (1981; 1982; 1983; 1985) and the racing model of Harris and Vickers (1987). These ideas have been applied to growth and macroeconomics (Aghion and Howitt, 1992; Aghion et al., 2001), market competition (Aghion et al., 2005) and antitrust (Segal and Whinston, 2007). As with the Chamberlin-inspired models, firms do not make a product choice in these models and, thus, differentiation between products is imposed exogenously.

Recent empirical work has demonstrated the simultaneous importance of horizontal and vertical differentiation. Braguinsky et al. (2021) provide evidence for this within firms and characterize rich paths of innovation, with discontinuous leaps in product characteristics followed by filling in of the newly created gaps. This matches our theoretical result of cycles

[5]Lancaster (1990) provides a thorough, albeit dated, review of the literature. Some models add incomplete or asymmetric information to the Hotelling formulation, although not in a way that captures innovation. For example, Meagher and Zauner (2004) incorporates uncertainty through a stochastic shock to demand that affects all products equally.

[6]Although the modern empirical literature relaxes this assumption, it does so only partially, retaining a degree of inflexibility in movement; see Arcidiacono et al. (2016).

between frontier and niche innovation at the market level.[7]

In building on the Hotelling framework, our model provides a sense of distance that captures the degree of novelty and riskiness of innovation. Letina (2016) and Bryan and Lemus (2017) develop models in which firms choose the direction of their innovation, although with a finite set of directions that correspond to different projects.[8] We allow for a continuum of correlated potential innovations. That uncertainty increases in the novelty of an innovation connects with Cabral (2003) in which a leader and a follower firm each make a binary choice of the variance of their innovation (high or low).

We use the Brownian motion to represent quality in a single dimension. This follows a recent search literature (Callander (2011); Garfagnini and Strulovici (2016); Callander and Matouschek (2019)).[9] The fundamental difference with that literature is that we allow for competition. Thus, firms take account of where other firms are located, whereas in the search literature the connection between agents is purely informational. This leads to fundamentally different insights into how competition and innovation interact that is not present in the search models. Callander and Matouschek (2022) study a static version of the model we analyze here, focusing on entry of a single firm and how innovation is affected by whether that entrant is independent or owned by an incumbent.

2 The Model

In every period $t = 1, 2, ...$ there is a continuum of consumers distributed uniformly on the product space $\mathcal{P} = \mathbb{R}_+$. A consumer $s \in \mathcal{P}$ who buys a product located at $l_j \in \mathcal{P}$ realizes gross utility,

$$u(s, l_j) = v(l_j) - \frac{1}{\tau}|s - l_j|, \tag{1}$$

where $v(l_j) \in \mathbb{R}$ denotes the quality of product l_j and $\tau > 0$ is an inverse measure of the degree of horizontal product differentiation. Net utility is obtained by deducting the price of product l_j from $u(s, l_j)$. In any given period, a consumer buys at most one unit of one product and consumes it immediately. The reservation utility of not consuming any product is zero.

[7] See Garcia-Macia, Hsieh and Klenow (2019) for economy-wide macro evidence on the relative importance of new product offerings versus improvements on existing products, although without distinction in the degree of novelty.

[8] The classic model of directed innovation is Acemoglu's (1998) work on labor versus capital-augmenting innovations. Our focus, in contrast, is on innovations within a single product market.

[9] The Brownian motion formulation resonates with the rugged landscapes literature in management, formalizing the idea that finding a good strategy is difficult (Levinthal, 1997; Rivkin, 2000). In that literature, search is blind, following variations on a hill-climbing algorithm rather than following optimal behavior based on well-formed beliefs, as it is here. That literature is also different in that it focuses on a search for organizational form within a firm, rather than the search for products in the face of market competition.

In period $t = 1$, a competitive fringe of firms supplies product $l_0 = 0$, which is known to have quality $v(l_0) > 0$. A firm can enter and develop a product at location $l_1 \in \mathcal{P}$. The competitive fringe implies that imitating an existing product cannot produce above normal profits. The new product is an experience good with expected quality $\mathrm{E}[v(l_1)|\, v(l_0)]$ and whose actual quality $v(l_1)$ is only revealed once it has been consumed.[10] We explain how expectations are formed below. Next, firms compete by setting prices simultaneously and independently. They are able to engage in third-degree price discrimination by charging different prices to consumers at different locations. Once firms have set their prices, each consumer decides what product to buy and consume. If at least some buy the new product, its quality is revealed publicly. Finally, firms realize their profits and time moves on to the next period.

Each subsequent period $t > 1$ proceeds analogously with the competitive fringe expanding to include the product introduced in period $t - 1$. The period t firm stays out of the market or enters at location l_t with expected quality $\mathrm{E}[v_t|\, \mathcal{E}_t]$, where $\mathcal{E}_t = \{l_0, l_1, ..., l_{t-1}\}$ is the set of existing products.

To the Hotelling framework we add uncertainty over the quality of innovations. To capture uncertainty, we represent the mapping $v(l_j)$ from product location to product quality as the realized path of a Brownian motion with zero drift and scale $\sigma > 0$. The firms do not know the path, and thus the quality of untried products. They do know the scale parameter σ and that the drift is zero, and they observe the quality produced by each product that is consumed. From the properties of the Brownian motion it follows that beliefs about the quality of a new product are normally distributed with a mean and variance that depend only on the known quality of the closest existing product in either direction.

For a new product that is beyond the frontier, beliefs depend only on the right-most existing product, which we refer to as the *frontier* product and denote by $l_t^f \equiv \max \mathcal{E}_t$. For a frontier innovation, $l_t > l_t^f$,

$$\mathrm{E}[v(l_t)|\, \mathcal{E}_t] = v(l_t^f) \tag{2}$$

and

$$\mathrm{Var}[v(l_t)|\, \mathcal{E}_t] = (l_t - l_t^f)\sigma^2. \tag{3}$$

The expected quality is the same as the frontier as drift is zero, and uncertainty is increasing in the distance from the frontier. This captures the intuition that uncertainty increases in the novelty of an innovation. In the first period, all new products are beyond the frontier at l_0.

Inside the frontier the existing products create a series of niches. Beliefs within each niche are a linear interpolation of the quality of the neighboring products in either direction.

[10]This assumption implies that the degree of uncertainty about true quality does not affect the profit from a new product in the first period when it is introduced. This aids considerably with tractability.

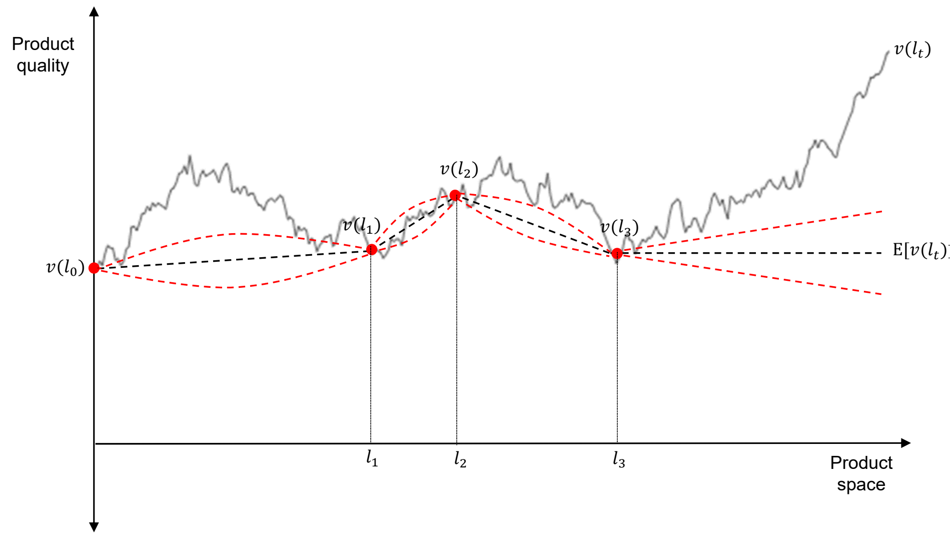

Figure 1: Beliefs on the Rugged Technological Landscape

Specifically, for any location l_t between neighboring products $l_L < l_R$,

$$\mathrm{E}\left[v\left(l_t\right)\right|\mathcal{E}_t] = v\left(l_L\right) + \frac{l_t - l_L}{l_R - l_L}\left(v\left(l_R\right) - v\left(l_L\right)\right) \tag{4}$$

and

$$\mathrm{Var}\left[v\left(l_t\right)\right|\mathcal{E}_t] = \frac{\left(l_t - l_L\right)\left(l_R - l_t\right)}{l_R - l_L}\sigma^2. \tag{5}$$

The variance of beliefs once again increases in novelty, reaching a peak in the center of the niche. The Brownian path and the beliefs it gives rise to when three new products have been introduced are depicted in Figure 1.

The development of new products is costly, requiring the investment of time and resources. These R&D costs typically increase in the novelty of an innovation along with the uncertainty about the outcome. Beyond the frontier, we suppose these costs are convex in novelty and, for simplicity, take the functional form:

$$c\left(l_t, \mathcal{E}_t\right) = \frac{c}{2}(l_t - l_t^f)^2,$$

with $c \geq \frac{2}{3\tau}$.[11] In a niche, new products are within the space that has already been researched, and development costs are not as high. For simplicity, we set this cost to zero, following the approach in Garfagnini and Strulovici (2016).

Entry and innovation continues in the market until it is no longer profitable to do so.

[11] This condition is sufficient to ensure the market is fully covered within its existing boundaries. Full covering is a standard assumption in the literature.

For simplicity, we set the production cost to be zero, such that the profit of an entering firm is its revenue less any R&D costs. We suppose additionally that the market ends with a small exogenous probability, $\gamma > 0$, in each period. This probability plays a role only in the simulations of Section 5.[12] The σ parameter scales the uncertainty in the market, and we refer to it as as the *complexity* of the technological landscape.

3 Developing New Products

3.1 Prices, Profits, and Competitive Shadows

The presence of a competitive fringe implies that profits are zero for existing products. In any period $t = 1, 2, ...$, therefore, the price of existing products is driven down to the cost of production, which we have assumed to be zero. For any consumer $s \in \mathcal{P}$, the best alternative to the new product l_t is to buy the existing product that maximizes gross utility (1), or to not buy a product at all. The value of this best alternative to the consumer is given by

$$f\left(s, \mathcal{E}_t\right) = \max\left\{0, u\left(s, l_0\right), u\left(s, l_1\right), ..., u\left(s, l_{t-1}\right)\right\}.$$

Third-degree price discrimination allows the entrant to set a price that extracts from each consumer all of the value it creates beyond this level. Specifically, given this best alternative, the highest, and profit-maximizing, price the entrant is able to charge consumer s in period t is given by

$$p\left(s, \mathcal{E}_t\right) = \max\left[0, \mathrm{E}\left[u\left(s, l_t\right)\middle| \mathcal{E}_t\right] - f\left(s, \mathcal{E}_t\right)\right],$$

where the consumer's expected gross utility from the new product is given by Equation (1).

The entering firm's profit in period t from location l_t is then given by:

$$\pi_t\left(l_t\middle| \mathcal{E}_t\right) = \int_0^\infty p\left(s, \mathcal{E}_t\right) \mathrm{d}s - c\left(l_t, \mathcal{E}_t\right).$$

The calculation of profit represents the classic dichotomy in business strategy between *value creation* and *value capture*. The value a product creates is $u\left(s, l_t\right)$ for the consumer at s. In a monopoly, this would equal the firm's profit, whereas in a competitive market the firm is not able to capture all of this value. Rather, the consumer captures $f\left(s, \mathcal{E}_t\right)$ and the remainder, $u\left(s, l_t\right) - f\left(s, \mathcal{E}_t\right)$, is what the firm captures as profit.

The decomposition of profit into value creation and value capture can be seen graphically in Figure 2. The value created by product l_2 is given by the triangle with a peak of $\mathrm{E}[v\left(l_2\right)]$

[12] See Footnote 16 for an explanation. γ may be thought of as the probability that an innovation in a neighboring technology space renders all products in this space obsolete.

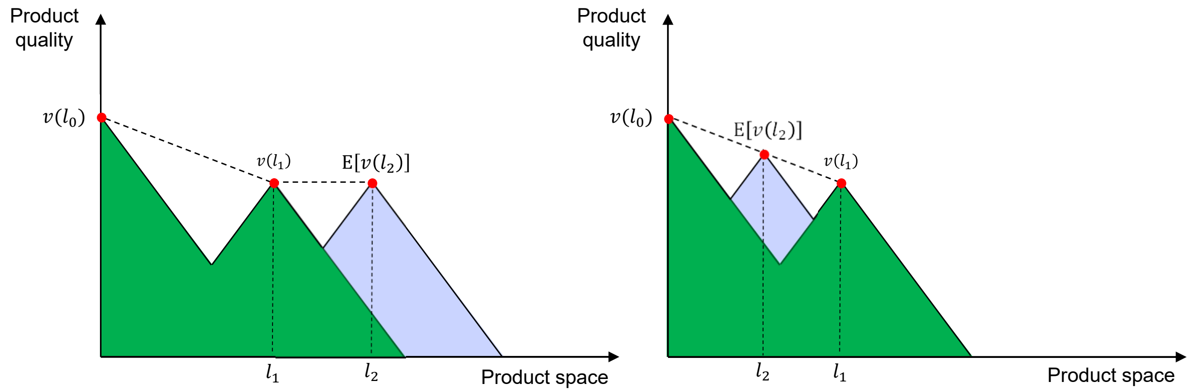

Figure 2: Value Creation & Value Capture on the Frontier (left) and in a Niche (right)

centered on l_2, and with sides of slope $\frac{1}{\tau}$. We refer to the triangle for each product as its *competitive shadow*. The profit of the entering firm—the value it is able to capture—is the part of its shadow that is above the competitive shadow of all other products. Any area that is also under the competitive shadow of another product is competed away and captured by consumers as consumer surplus. Each panel of Figure 2 depicts a potential entrant at l_2: On the frontier in the left panel, and in the niche between products l_0 and l_1 in the right panel. As these are experience goods, the height of the new product's competitive shadow in the period of entry is given by the expected quality of its innovation. The profit of each entrant is then the blue region.

3.2 Frontier Innovation

In the first period, the firm must innovate on the frontier if it innovates at all, and the only question is how far to the right of l_0 its product should be located. The size of the firm's competitive shadow is the same wherever it locates on the frontier. What changes is how much of that shadow is above the shadow of the competitive fringe at l_0. (This logic is evident in the left panel of Figure 2 even though it depicts a later period.) As the entrant locates further to the right, it captures more of the value that it creates, thereby increasing its gross profit. From this must be subtracted R&D costs, which increase the further to the right the firm locates. The optimal choice is given by Proposition 1.

Proposition 1 *In period one, the entrant innovates and its optimal location is*

$$l_1^* = \frac{2\tau v\left(l_0\right)}{1+2c\tau}$$

and profit is

$$\pi_1 = \tau v (l_0)^2 - \frac{2c\tau^2}{(1+2c\tau)} v (l_0)^2 .$$

Frontier-expanding innovation involves a combination of competing for existing customers and bringing new customers into the market. The more the entrant differentiates, the more it focuses on new customers and the more it softens competition for existing customers. This combination implies that differentiation strictly increases the value that is captured by the firm. Value creation and value capture are not in conflict in this situation, and innovation is constrained only by R&D costs.

Innovation at the frontier increases in the quality of the incumbent product; l_1^* is strictly increasing in $v(l_0)$. Thus, successful products themselves induce bolder innovation. One intuition may be that a higher quality incumbent is a more fearsome competitor and this incentivizes the entrant to differentiate itself more. Missing from this intuition is that the entrant itself is of higher quality, in fact equally so. The reason the entrant differentiates more is that higher quality products create more value and, thus, there is more to lose to consumer surplus by competing. Geometrically, more of the entrant's competitive shadow emerges from the incumbent's shadow the higher are the product qualities, and, thus, the more the entrant differentiates in equilibrium.

After the first period, the frontier moves to l_1 and a niche opens up between l_1 and l_0. The second firm faces a logic at the frontier that is similar but different to that of the first firm. The difference that can arise is that the frontier product may not be *active* in the market. A product is active if it attracts customers and, thus, some part of its competitive shadow is above all others. Denote the set of active products $\mathcal{A}_t$, with the largest active product given by $l_t^a \equiv \max \mathcal{A}_t$. If the frontier product is active then $l_t^a = l_t^f$, otherwise $l_t^a < l_t^f$.

If the frontier product is active, then it is the relevant competitor for any new product on the frontier. This was the situation in the first period. If the same is true in later periods, then the logic of the first period carries over directly to later periods with product l_0 replaced by the current frontier product, l_t^f.

The logic is different if the frontier product is inactive. An inactive frontier product means the entrant's competitive and technological opportunities are separated. The entrant must compete against a product that is inside the frontier, but its technological prospects are dictated by an inferior product at the frontier. This is depicted in Figure 3.

An inactive frontier product makes innovation more difficult and potentially stifles it altogether. If the frontier product is too deeply embedded within the competitive shadow of another product, an entrant must experiment boldly to simply escape the shadow and find a product that will be active. How deeply embedded the frontier can be before frontier innovation is no longer profitable is a relative rather than an absolute standard. Specifically,

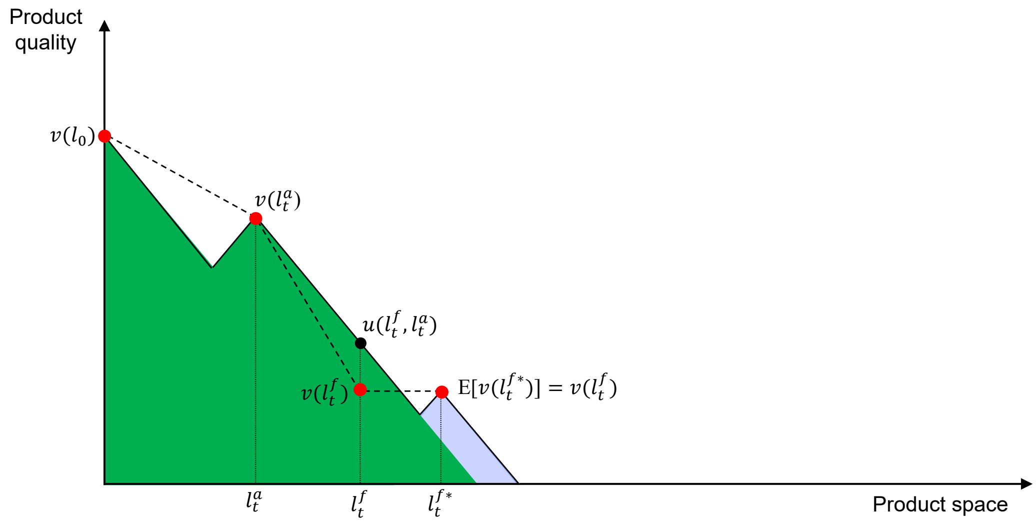

Figure 3: Frontier-Expanding Innovation with a Non-Active Frontier Product

if the quality of the frontier product exceeds a fixed fraction of the shadow in which it is embedded, then frontier innovation remains profitable. The height of the shadow in which the frontier is embedded is equal to $u(l_t^f, l_t^a)$, the gross utility of the largest active product for the consumer at the frontier product's location. We then have the following.

Proposition 2 *There is a constant $\kappa \in (0,1)$ such that, in any period $t \geq 2$,*
(i) if $v\left(l_t^f\right) \leq \kappa \cdot u\left(l_t^f, l_t^a\right)$, there is no profitable frontier innovation.
(ii) if $v\left(l_t^f\right) > \kappa \cdot u\left(l_t^f, l_t^a\right)$, the optimal frontier innovation is located at

$$l_t^{f*} = l_t^f + \frac{\tau}{1+2c\tau}\left(v\left(l_t^f\right) + u\left(l_t^f, l_t^a\right)\right)$$

and generates profit

$$\pi_t\left(l_t^{f*}\right) = \tau v\left(l_t^f\right)^2 - \frac{c\tau^2}{2\left(1+2c\tau\right)}\left(v\left(l_t^f\right) + u\left(l_t^f, l_t^a\right)^2\right) > 0.$$

If the frontier product is active, then $u(l_t^f, l_t^a) = v(l_t^f)$ and the expression for l_t^{f*} is equivalent to that in Proposition 1. If the frontier product is inactive, the entrant's optimal innovation depends on the quality of both the frontier and the largest active products, and is increasing in both.

The profit from innovation is also increasing in the quality of the frontier product but it is decreasing in the quality of the largest active product when these are different. This can be seen in the profit function in Proposition 2 where the quality of the largest active product is proxied by $u(l_t^f, l_t^a)$. The ratio κ implies that the higher the quality of the frontier and

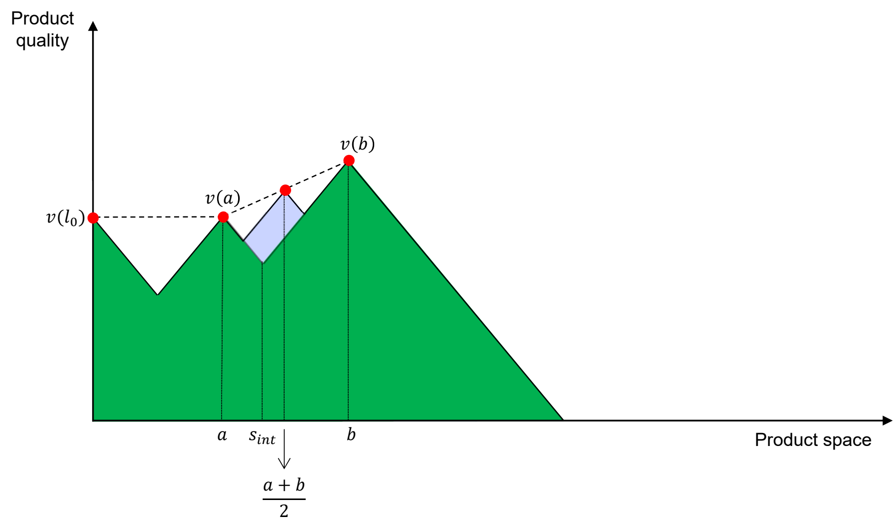

Figure 4: Niche Innovation

active products are, the more deeply embedded in an absolute sense can the frontier product be before frontier innovation is no longer profitable, reflecting the better profit opportunities when frontier quality is higher.[13]

Proposition 2 establishes that frontier innovation can end and that it can end inefficiently early, even when the frontier is expected to produce positive quality products. Although it may not be in the interests of a single firm to innovate, it is still socially efficient to do so as the firm's calculation ignores the ongoing consumer surplus from the new product, not to mention that the firm's innovation may induce valuable follow-on innovation. Indeed, these benefits may be so high that continuing innovation is socially beneficial even when the quality at the frontier is negative. The premature end of innovation represents a market failure, one that we return to in Section 6.

3.3 Niche Innovation

Innovation in a niche does not offer the possibility for an entrant to escape competition. The question for the entrant, then, is not how to avoid competition but rather who to compete with.

The answer is not straightforward. Consider the situation depicted in Figure 4. In the niche between products a and b, the entrant faces contrasting competitors. To the right, b is a high quality product, whereas to the left is the lower quality product, a. It is natural to think that the entrant would gravitate toward the product on the left as it is weaker and less of a

[13]That profitable opportunities exhaust themselves according to the fixed ratio κ can be seen by setting the profit function in Proposition 2 equal to zero. This gives a closed-form expression for κ, which we provide in the appendix.

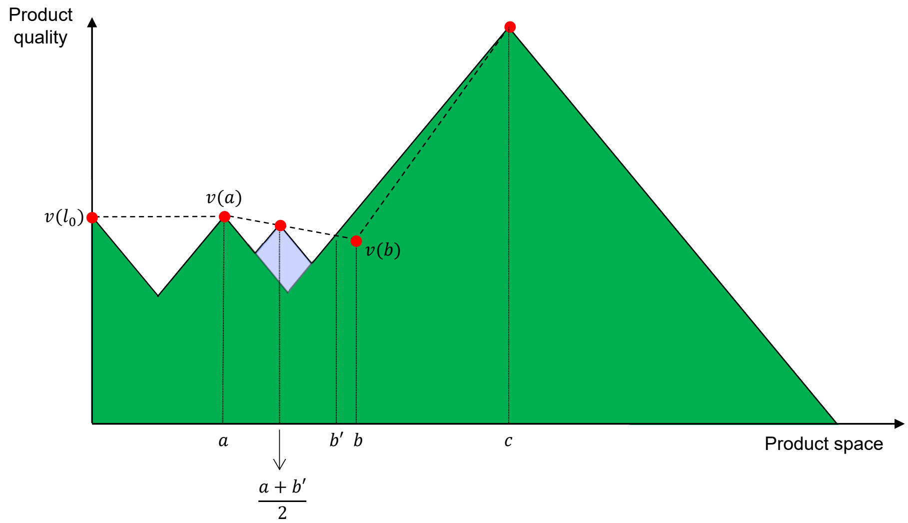

Figure 5: Niche Innovation on the Viable Sub-niche $[a, b']$

competitive threat. In so doing, however, the entrant's expectations about its own product are weakening. The other end of the niche offers the opposite trade-off. The competition is of higher quality, but so too is the expected quality of the entrant's own product.

This creates a dilemma. The entrant must choose between a strong product and a weak competitor, it cannot have both. This calculus is complicated further when one, or both, of the ends of a niche are inactive, such as for the niche between products a and b in Figure 5. In this case, moving toward the weaker neighbor lowers the expected quality of the entrant's product but leads it closer to a stronger competitor. If it gets too close to the inactive neighbor, the new product will itself be inactive.

Despite the richness of possibilities, the optimal entry strategy for a firm in a niche takes on a simple form. A firm enters only in the part of the niche that is above the competitive shadow of all competing firms and, within that portion, it locates exactly halfway along, regardless of the width of the niche or the quality of the neighboring products.

We define the part of the niche that is above the competitive shadows of other products as a *viable* niche. In Figure 4, the entire niche $[a, b]$ is viable, whereas in Figure 5, the viable niche is the subset $[a, b']$. In some cases, a niche contains no viable section, such as for niche $[b, c]$ in Figure 5. We then have the following result.

Proposition 3 *Suppose firm t locates in the viable niche $[a, b]$. Its optimal location is then given by*

$$l_t^{n*}(a, b) = \frac{1}{2}(a + b)$$

and its profits are given by

$$\pi_t\left(l_t^{n*}(a,b)\right) = \frac{1}{8\tau}(b-a)^2\left(1-\tau^2\beta(a,b)^2\right),$$

where $\beta(a,b)$ *is the slope of expected quality in the niche. Innovation in a non-viable niche is not profitable.*

For a niche that is entirely viable, as in Figure 4, the entrant's location is where uncertainty about the innovation is maximized and competitive differentiation is at its largest. Proposition 3 shows, however, that these properties are incidental. This is evident when the viable niche is a subset of the total niche. Then, in locating halfway along the viable niche, the entrant is closer to one neighbor than the other, and, in fact, closer to the competitive neighbor if only one is active. Moreover, uncertainty about the innovation at this choice is below the maximum that is possible.

This choice reflects a trade-off between market power and market share. The entrant's market power with any consumer is maximized when it locates at the point where the competitive shadows of its neighbors intersect, marked as s_{int} in Figure 4. The consumer at this point is receiving the lowest utility of consumers in the niche and, thus, has the highest willingness to pay should the new product be targeted directly to its preferences.

In a *balanced* niche—where the neighbors are of equal quality—this location is exactly halfway along the niche and maximally differentiates the entrant from its competitors, thus minimizing competition. In an *unbalanced* niche—with neighbors of unequal quality—moving toward the higher quality neighbor increases both the value created by the entrant and its market share. However, it does so at the cost of more intense competition. Halfway along the viable part of a niche is where the pursuit of market power and market share optimally balance out.

Even though the optimal location in a niche is independent of the width, height, and slope of the niche, the profit the entrant earns does depend on these properties. Inspecting the profit function in Proposition 3 reveals that the entrant's profit is independent of the absolute quality of its neighbors and, indeed, the expected quality of its own product. This is evident in Figure 4. As the entrant captures the part of its competitive shadow that is above those of its active neighbors, if those neighbors increase or decrease in quality without the slope of the niche changing, the profit of the entrant is unchanged.

The entrant's profit does depend on the width and slope of the niche. As the slope increases—and the niche grows more unbalanced—the entrant's ability to capture the value it creates decreases, whereas the value it creates increases at a lower rate (or even decreases). Thus, the entrant prefers to compete against neighbors who are evenly matched in a balanced niche rather than face one stronger and one weaker competitor.

This implies that a firm will choose a niche not based on the expected quality of the innovation—on the value that it will create—but on the relative nature of competition it will face from its neighbors. In particular, it will choose a niche that is balanced, even if that niche is narrower, and even if the expected quality of its product is lower. Thus, a firm will sacrifice its ability to horizontally differentiate from its competitors, and sacrifice its own expected quality, if in so doing it finds a more balanced and thus hospitable competitive environment.

This is a striking implication as it says that firms will deliberately enter below the quality frontier. This is not the logic of an entrant positioning at the low end of the market to differentiate and avoid competition. Rather, this is an anti-differentiation result. The entrant locates below the quality frontier precisely because it expects to be of similar quality to its competitors.

4 The Dynamics of Innovation and Industry Structure

The behavior of individual firms aggregate into the dynamics of innovation and industry structure. The rugged technological landscape generates a dynamic path that is rich and irregular, exhibiting a wide variety of market structures and innovation dynamics. A complete characterization of these dynamics is not possible. We present a partial characterization, focusing on disruption, switches in the type of innovation, and when innovation stops.

Frontier Innovation. If an entrant innovates at the frontier, what type of innovation comes next? Does the next entrant continue frontier innovation? Or does it reverse course and pursue a niche? Proposition 4 addressees these questions. Recall that $\kappa \cdot u(l_t^f, l_t^a)$ is the threshold in Proposition 2 above which frontier innovation is profitable, and $u(l_t^f, l_t^a)$ is the level at which the frontier product is itself active.

Proposition 4 *Suppose that frontier innovation is optimal for firm $t-1 \geq 1$. Then there exists a $\overline{v}_t^f \in (\kappa \cdot u(l_t^f, l_t^a), u(l_t^f, l_t^a))$ such that:*
(i) if $v(l_t^f) > \overline{v}_t^f$, frontier innovation is optimal for firm t.
(ii) if $v(l_t^f) \in (\kappa \cdot u(l_t^f, l_t^a), \overline{v}_t^f]$, frontier innovation is not optimal for firm t, but may be optimal for some firm $t' > t$.
(iii) if $v(l_t^f) \leq \kappa \cdot u(l_t^f, l_t^a)$, frontier innovation is not optimal for firm t or any firm $t' > t$.

Case (i) follows from the logic of Proposition 2, albeit with a subtlety. Success at the frontier makes the frontier more attractive. Thus, if the frontier dominates all niches for firm t, a successful outcome, or even an outcome that is not much worse, means that the frontier again dominates the same niches. The subtlety is that the frontier innovation by firm t itself creates an additional niche. If this niche is balanced, it may be more attractive than previous niches. Proposition 2 establishes that, even when this occurs, it does not dominate the frontier, and a single threshold determines when frontier innovation continues. As beliefs are normally distributed and the threshold for continuation is below the previous level, the probability that frontier innovation is followed by further frontier innovation is strictly greater than 50%.

For cases (ii) and (iii) frontier innovation disappoints and firms turn away from the frontier. In case (iii) the frontier performance is so bad that frontier innovation ends forever. In case (ii) frontier innovation remains profitable for now, though it is not guaranteed that it will recommence. As we will see below, a niche innovation may be so successful that it disrupts the frontier product, closing off frontier (and possibly all) innovation thereafter.

Frontier innovation can also have an impact on niche innovation. If a frontier innovation is a breakthrough success, it not only opens a door to further frontier innovation, it closes the door on niche innovation, at least in the parts of the product space explored so far. In this sense, success at the frontier brings consumers into the market and it also disrupts existing products, winning their customers and driving those products from the market. If the frontier innovation is of sufficiently high quality it disrupts the entire market. Moreover, when it does so, it turns the parts of the product space that have already been explored into a "dead zone" in which no future firm will ever locate. The entrant then obtains a competitive "moat" due to its own success and the information gleaned from the lower quality of its predecessors.

Corollary 1 *There exists a* $\overline{\overline{v}}_t^f > v(l_t^a)$ *such that if* $v(l_t) > \overline{\overline{v}}_t^f$ *no entrant ever locates to the left of* l_t *again.*

As devastating as frontier disruption can be to innovation, it is only one-sided. Expectations on the frontier are increased by frontier success and frontier innovation becomes even more attractive. Indeed, niche innovation stops only to the left, as the frontier innovation that follows creates new niches that can then be exploited. In fact, the disruptive innovation itself soon becomes part of a niche and may be disrupted by a future breakthrough innovation.

Niche Innovation. Niche innovation follows a different logic to that at the frontier. Rather than opening a door, successful niche innovation closes the door on further innovation. This does not imply a complete inversion of the logic of frontier innovation. Failure in a niche also deters future innovation. Instead, it is middling performance that allows further innovation, whereas extreme performance in either direction shuts it down.

Proposition 5 *Suppose that firm* t *innovates in niche* $[a, b]$. *There are thresholds* $\underline{v}_t^n(a, b) < \overline{v}_t^n(a, b)$ *such that if* $v(l_t) \in (\underline{v}_t^n(a, b), \overline{v}_t^n(a, b))$ *then it is profitable for firm* $t+1$ *to locate within* $[a, b]$. *If* $v(l_t) \notin (\underline{v}_t^n(a, b), \overline{v}_t^n(a, b))$ *then no firm firm locates within* $[a, b]$ *again.*

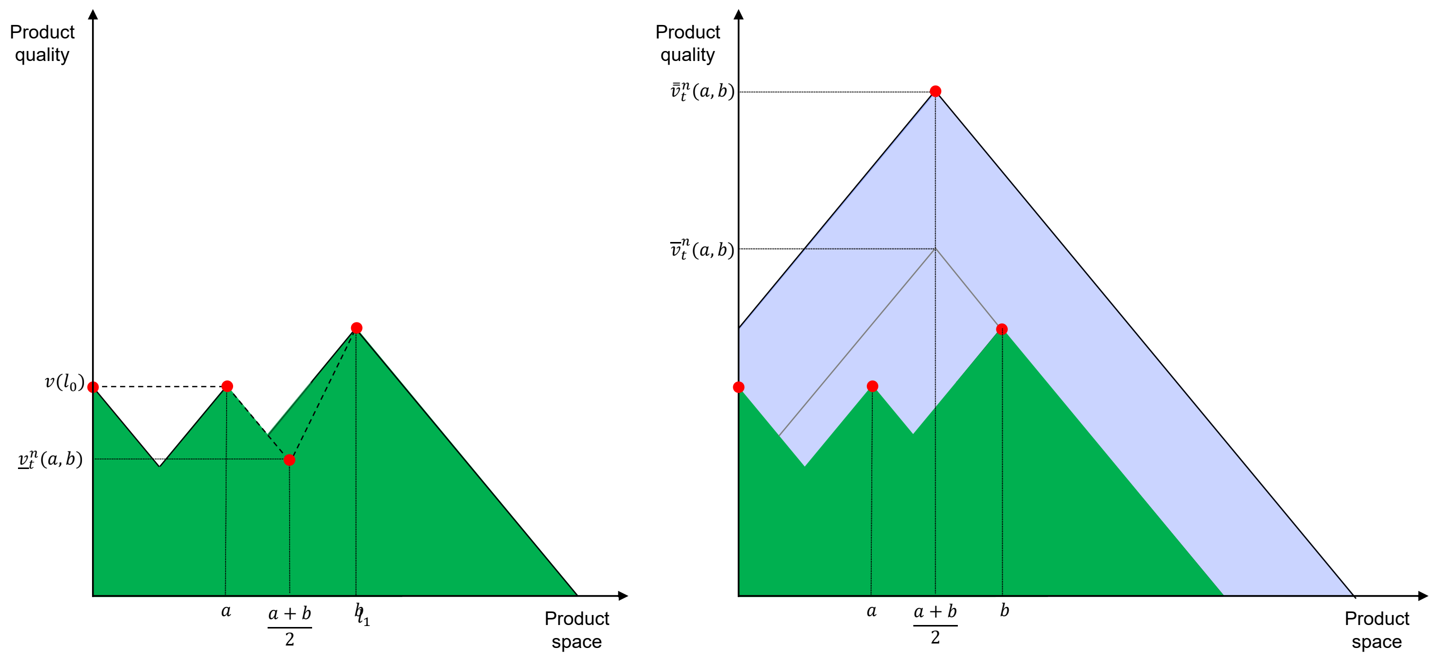

Figure 6: Entry Deterring Niche Innovation from a Failure (left) and a Success (right)

Innovation in a niche splits the niche into two. This leaves less room for differentiation, which by itself makes further innovation within the niche less likely. Moreover, if the realized outcome is high or low, the two niches created will be unbalanced and leave no profitable opportunities to exploit. The two thresholds are depicted in Figure 6. In either case, all remaining products in the original niche lie in the competitive shadow of another product. It is only if the newly created niches are relatively balanced that profitable opportunities remain.

Though success and failure in a niche both deter future innovation within the niche itself, they have very different impacts on the broader market. Failure in a niche is contained within the niche, and has no impact beyond that. In contrast, successful innovation in a niche can deter innovation beyond the niche itself. This is evident in Figure 6 as should the realized outcome be sufficiently high, the new product overshadows not only the niche between l_L and l_R, but also all neighboring niches as well as the frontier.

In this way, successful niche innovation closes the door to further innovation. In fact, for sufficiently large breakthroughs, niche innovation can disrupt the entire market, driving all existing products—in both directions—from the market and shutting down innovation for ever more.

Corollary 2 *There exists a* $\overline{\overline{v}}_t^n(a,b) > \overline{v}_t^n(a,b)$ *such that* $v(l_t) \geq \overline{\overline{v}}_t^n(a,b)$ *implies innovation stops forever.*

Niche disruption is more damaging to innovation than is frontier disruption as it is two-sided. The one-sided disruption of frontier innovation offered the silver lining that frontier innovation became more attractive. No such silver lining exists for two-sided disruption of niche innovation.[14]

The End of Innovation. Corollary 2 shows that it is possible for innovation to end and demonstrates one way in which it can happen. A second way in which innovation can end is slower and less dramatic. A failure at the frontier closes off the frontier, bounding innovation thereafter but without necessarily stopping it. New entrants may turn to niche innovation, seeking out opportunities within the existing market. This can last a long time, yet it cannot last forever. Moreover, it need not end with disruption nor with one product dominating the market. Rather, it may simply just exhaust itself and peter out.

Proposition 6 establishes this result and shows generally that innovation cannot continue indefinitely in a bounded region of the product space. This holds despite the fact that R&D costs in a niche are zero.

Proposition 6 *Given l_t, innovation in the interval $[0, l_t]$ contains at most a finite number of products almost surely.*

It follows, therefore, that for innovation to have an engine of growth, it must come from the frontier. If frontier innovation is no longer profitable, not only will the market stop growing in size, but the gains from innovation are thereafter bounded and inevitably will come to an end.

These properties contrast with entry in Hotelling's classic model. With constant product quality and consumers arrayed across the real line, entry and market growth continue indefinitely on the frontier. Even if the product space were bounded and innovation forced to be in niches, cost-free niche entry would continue in perpetuity. In a world of innovation, in contrast, growth is inspired by the ruggedness of the technological landscape and it is also constrained by it. Uncertainty over the locations of the peaks and valleys can undermine the incentive to innovate.

The long-term prospects at the frontier depend on the drift term. If it is even slightly negative, it is trivial that innovation must eventually end with probability one. If the drift is positive, there exists an escape probability such that quality is so high, and each innovation is sufficiently novel, that the probability innovation ends approaches zero. The answer for the case of zero drift is not obvious, and as zero serves only as a neutral benchmark, we do not investigate this question further here.[15]

[14] This suggests that forward-looking firms would value niche disruption even more than frontier innovation.

[15] The qualitative properties of search that we focus on are not affected by whether the drift is on one side of zero or the other.

5 Numerical Analysis of Market Dynamics

5.1 A Tale of Three Markets

To see the richness and variety of market dynamics that are possible, it is helpful to begin with three markets that exhibit contrasting patterns of innovation. Figure 7 depicts a market in which there is little innovation and that stabilizes quickly. As can be seen in the left panel, Only four products are introduced before the market stabilizes, three frontier innovations followed by a niche innovation. Of these, only one, l_4, along with product l_0, remain active when stability is reached.

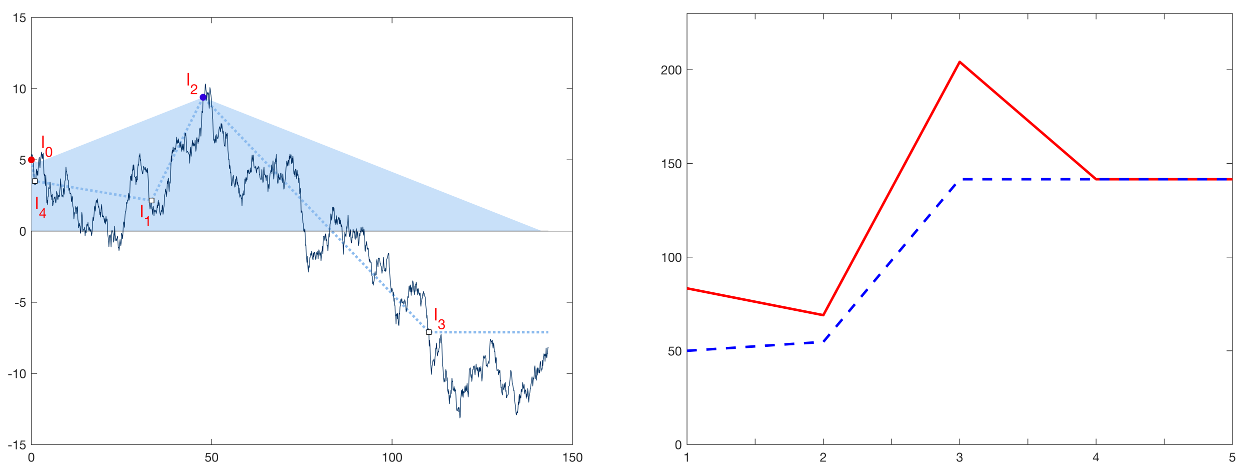

Figure 7: Short-Lived Innovation: New Products (left) and Market Size Growth (right)

The trajectory of market size is shown by the red line in the right panel. The blue line is the measure of consumers served by the incumbent products (i.e., the competitive fringe). The blue line grows monotonically, therefore, as new products are introduced and settle into their place in the market. The difference between the red and blue lines is the consumers brought into the market by the new product. The red line need not be monotonic as new consumers who are drawn in by the promise of the new product leave the market if the product underperforms and is not followed up with a new product nearby. For the market in Figure 7 the market size falls back to the blue line upon the failure of product l_3, and it remains there as the subsequent niche innovation does not bring new consumers into the market.

The two markets depicted in Figure 8 are very different. In both of these markets innovation is extensive and long-lasting with hundreds of new products introduced. The growth trajectory in each of these markets is very different, however, with each generating a different pattern of innovation. In the left market, innovation is largely in two phases. The market begins with a long phase of frontier innovation followed by niche innovation as new entrants exploit the many niches that were opened up. In contrast, the market on the right exhibits

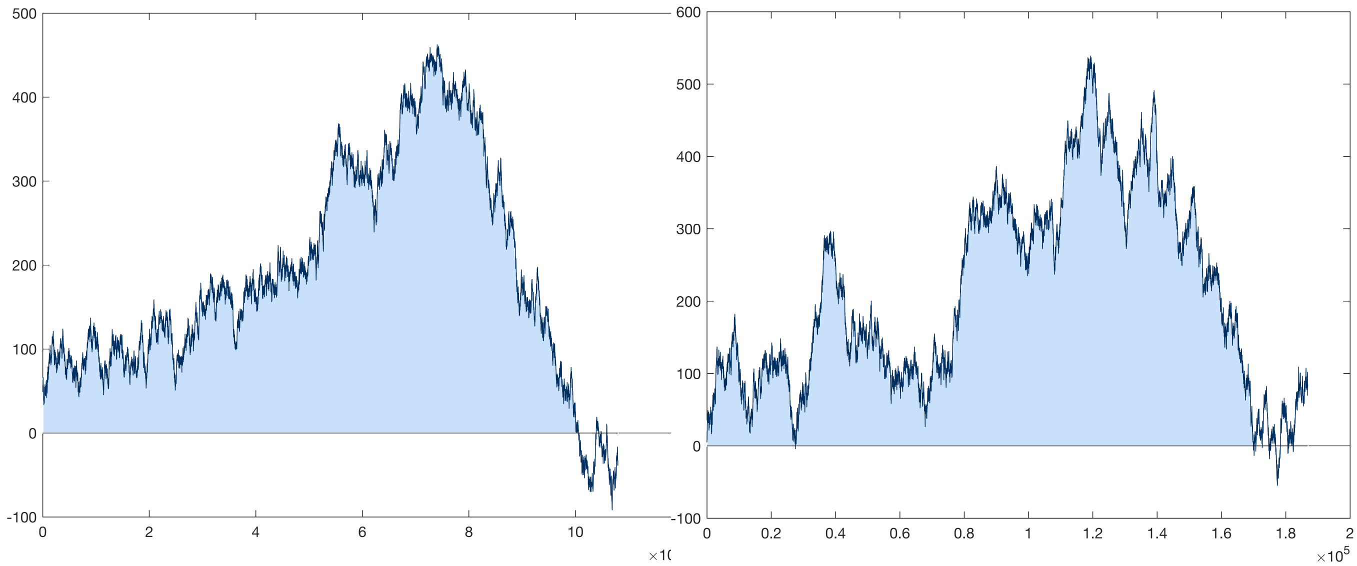

Figure 8: Long-Lived Innovation

more changes in direction and it cycles between frontier and niche innovation. As in the left-side market, innovation proceeds in a broad sweep that expands the frontier followed by niche innovation, although in this market there are three such eras and each is shorter relative to the left-side market. (In both markets the separation between frontier and niche innovation is not sharp, with the odd period of niche innovation appearing during a phase of frontier innovation, and vice versa.)

The different patterns in innovation can be seen clearly in the trajectory of market size. These are depicted in Figure 9 (the red and blue lines are indistinguishable in these markets). The left-side market grows in one continuous arc before settling off and remaining relatively stable thereafter as niche innovation plays out. In contrast, the right-side market has three distinct growth phases. Each increase represents an era of frontier-expanding innovation with each plateau reflecting the subsequent niche innovation.

The innovation pattern in a market is tightly linked to the technology landscape. The left-side market resembles one broad mountain with many subpeaks. In contrast, the right-side market is more aptly described as several distinct mountains. A useful heuristic is the notion of *prominence* from topography. Topographic prominence measures the height of a mountain's summit relative to the lowest contour line encircling it but containing no higher summit within it. We know also Proposition 2 that the profitability of frontier innovation depends on how rapidly the frontier height falls from the previous peak. Combining these ideas, we can see that the greater the prominence of a peak in the technological landscape, and the faster it falls from its peak, the more likely is a market to turn from the frontier to niche innovation. In the left-side market, the peaks are not as prominent, and the fall not so rapid, that frontier innovation continues until both of these conditions fail, and spectacularly

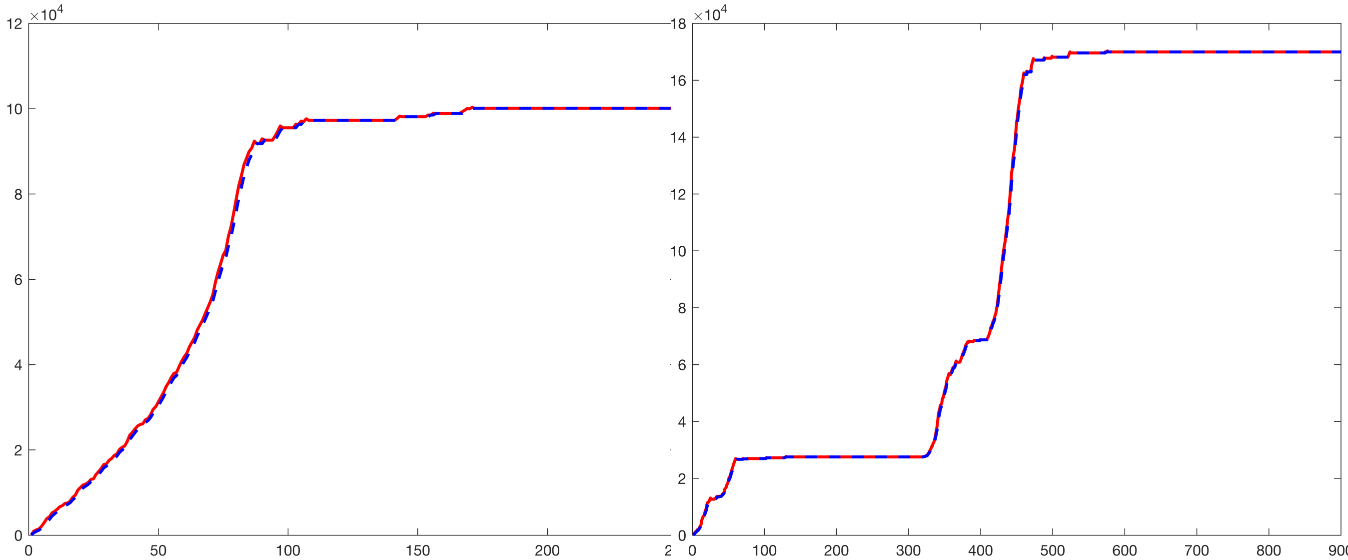

Figure 9: Market Size Growth

so, that frontier innovation permanently ends. The different trajectory of innovation in the right-side market follows from the greater prominence of the early peaks, and the sharp falls to close to sea level, so to speak.

5.2 Averages and Comparative Statics

These markets represent several possibilities among a broad array of market structures. Each of the three markets were generated from a simulation with parameters $v_0 = 5$, $\sigma = 1$, $\tau = 10$, $c = 0.1$, and $\gamma = 0.001$. We report here averages for this benchmark set of parameters across 200,000 simulations.

The most prominent feature of market dynamics in the simulations is a long right tail. The average number of products introduced is 49.3, although the range extends up to 9,384 products.[16] Firms innovate at the frontier 43.9% of the time and in a niche 56.1% of the time. The maximum number of cycles between frontier to niche innovation is 38, with an average of 1.6. Thus, many markets transition only once or a few times between innovation types, although the tail is long and some markets work through many cycles. In 95.1% of markets innovation stopped endogenously when profit opportunities were exhausted; in the remaining 4.9% of markets innovation stopped exogenously (due to the γ parameter).

To yield further insight into the economic forces within the model, we turn to comparative statics. Table 1 reports summary statistics for variation in the degree of product substitutabil-

[16]It is for this reason that we allow the market to end exogenously in each period ($\gamma = 0.001$). Without this possibility, some markets would continue indefinitely, not only exceeding the limits of computation but also skewing our results. This strikes us as more reasonable than truncating the product space at some arbitrary value.

ity (τ), R&D cost (c), and the complexity of the technological landscape (σ).[17]

Product substitutability: The τ parameter speaks directly to the classic question of the relationship between competition and innovation. Increasing τ increases competitive intensity as consumers are more easily able to substitute products for each other. In our model an increase in τ has a dampening effect on innovation. An increase in τ from 10 to 20 decreases the number of new products introduced from 49.3 to 14.9.

To understand why competition and innovation are negatively related, we need to break down the type of innovation that is undertaken. The direct effect of an increase in τ is felt within the market boundaries. Profitable opportunities for niche innovation are fewer and less lucrative, and this deters innovation. Along with this is an indirect effect beyond the market boundaries. With niche innovation less attractive, firms switch to frontier innovation, and the fraction of frontier innovations increases from 43.9% to 54.7%. Moreover, the firms innovate more boldly at the frontier, differentiating themselves further from existing products to reduce competition. Consequently, they leave behind larger niches that are more profitable to exploit. Nevertheless, the direct effect dominates the indirect effect and increased competitive intensity decreases the total number of innovations.

Despite a dramatic decrease in the number of products, an increase in τ increases welfare by 35%.[18] The change in welfare also reflects direct and indirect effects. For a given set of products, an increase in τ increases welfare directly as consumers now experience greater utility from those products. The indirect effect is the effect on innovation, with fewer new products when τ is higher. Although the reduction in products is dramatic, it is not so great as to negate the direct effect, and welfare increases.

R&D costs: A change in R&D costs also yields direct and indirect effects. As R&D costs apply only on the frontier, the direct effect is that higher costs push firms away from the frontier and toward niche innovation. Yet niches are eventually exhausted and frontier innovation does continue. When it does, it is less bold. As a result, the niches left behind are narrower and less attractive, and this renders frontier innovation more attractive. We find, surprisingly, that this indirect effect dominates the direct effect. As a result, an increase in R&D costs actually leads to *more* frontier innovation, not only relative to niche innovation but in absolute terms too.

This finding provides a novel twist on on how innovation is interpreted in practice. A market with high R&D costs has a higher number and higher proportion of frontier-expanding innovations and, therefore, may appear highly innovative. However, this is misleading. If

[17]We explored many parameter values for all variables and the substantive conclusions report do not seem to depend on the values used.

[18]Welfare is the sum of consumer and producer surplus; i.e., the area under all competitive shadows.

	Benchmark	$\tau = 20$	$c = 0.2$	$\sigma = 2.0$
New products	49.3	14.9	30.0	19.6
Frontier products %	43.9%	54.7%	50.9%	52.2%
Welfare	12,078	16,337	4,316	59,035

Table 1: Comparative statics: $v_0 = 5, \sigma = 1, \tau = 10, c = 0.1,\ \gamma = 0.001,\ \delta = 0.9,$ 200,000 simulations.

one were to look at the scale of innovation, it would become clear that each innovation is more incremental. In our simulations, the higher propensity of frontier innovation actually represents a decrease in welfare of 64%.

Complexity of the Technological Landscape: An increase in σ increases the ruggedness of the technological landscape. A classic insight from strategic experimentation is that higher variance is good for experimentation and innovation. In our model, this insight applies at the aggregate level but not at the level of individual firms. A more rugged landscape deters innovation within niches, following the logic of Proposition 5. As a result, a higher proportion of innovations are at the frontier, yet even there the chance of an innovation-stopping failure is higher. The average number of new products falls from 49.3 to 19.6 when σ increases from 1 to 2. However, the innovations that are undertaken are more valuable. An increase in the ruggedness of the landscape means there are more failed products but also more breakthrough successes. These breakthroughs are of sufficient quality that, despite the dampening effect they have on subsequent innovation, total welfare increases by close to 500%.

5.3 The Life Cycle of Products and Firms

Embedded within the dynamics of innovation lies a divide between those products that succeed and those that fail to find a place in the market. Even among those that fail, some fail immediately upon entry, whereas others initially survive, and even thrive, before being disrupted by a later entrant.[19] The ratio between these outcomes appear surprisingly stable across many permutations of parameters. For the benchmark parameters, an average of 37.0% of products remain viable when the market stabilizes, and for those that disappear the average life-span in the market is 14.1 periods for niche innovations and 22.3 periods for frontier innovations.[20]

Despite the longer expected life-span for frontier innovations, they are less likely to survive

[19] We emphasize *products* here rather than *firms*, as our assumption on the competitive fringe implies no firm makes an above normal profit after its period of entry.

[20] Recall that the distribution in the number of products has a long tail. When many products are introduced there are more products driven from the market and this occurs over a longer time frame. This generates the relatively long average lifespan for products that eventually leave the market.

through to market maturity. Frontier innovations represent 43.9% of total innovations, yet they represent only 24.5% of products that are active when the market stabilizes. This finding runs counter to the ideal of a bold innovator reaping the benefits of her breakthrough. It resonates instead with the perception that bold innovators capture little of the value they create.

The reason for the divergence between life-span and long-run success can be found in the degree to which innovation begets further innovation. At the frontier, success encourages further innovation, with later firms "standing on the shoulders of giants." This complementarity is good for society and is welfare enhancing, yet it is bad for the initial innovator as competition will become more intense. In opening a door for others to follow, a frontier innovator is condemning itself to likely obsolescence.

In contrast, a successful niche innovation closes the door to further innovation. In not providing shoulders for others to stand on, a niche innovation is insulated better for the long term. This is detrimental to society, but more valuable to the innovator itself. This implies that far-sighted firms may deliberately seek niche rather than frontier innovations, despite the fact that doing so will lower the overall societal benefit of innovation.

6 Policy Interventions

In many markets, innovation stops inefficiently early. A breakthrough success in a niche or a failure at the frontier can put an end to innovation, even with much of the product space still unexplored and many consumers unserved.

For innovation at the frontier to stop, a failure need not be absolute in the sense that the product creates no positive value. All that is required is that a frontier innovation falls sufficiently far into the competitive shadow of another product, even if its quality is positive.

The inefficiency that results can be large. Figure 9 depicts a market in which innovation stops inefficiently early due a positive but low quality for product l_2. As can be seen, the value captured in the market represents only a small fraction of what is available. To an outsider, this market would look like a mild success or even a disappointment. A few products were introduced that expanded the scope of the market and served more customers, but without improving on the initial product l_0. Unbeknownst to all is that the technology is highly promising, and that the failure of l_2 was only a blip, albeit a particularly unlucky one.

Although innovation is not profitable for an individual firm, it may still be valuable to society, both because an innovation creates consumer surplus that is not factored in by the firm, and because a successful innovation can spur further innovation. This represents a market failure and it suggests an opportunity for productive government intervention. If a policy can incentivize innovation that crosses a valley in the landscape, such as the one l_2

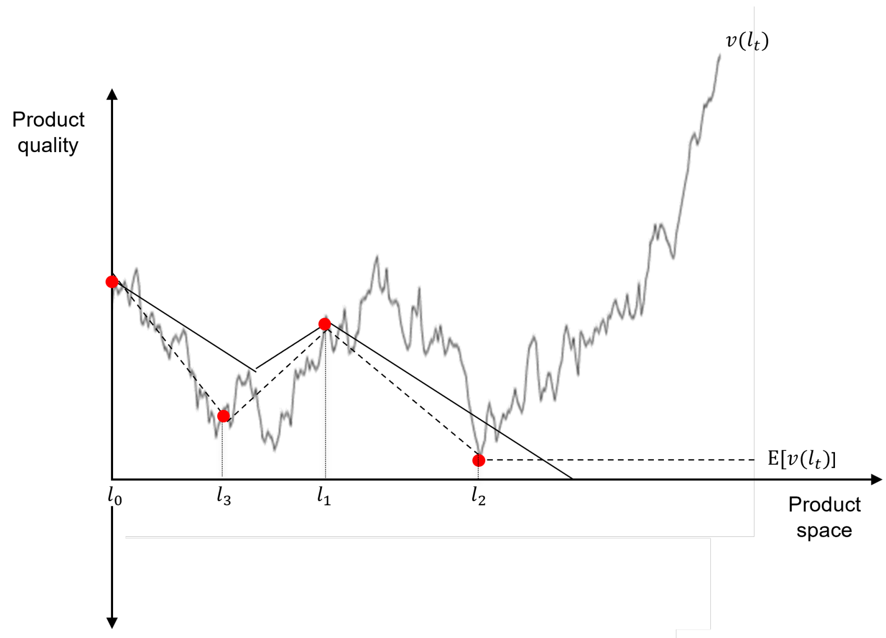

Figure 9: Incomplete Innovation

sits within, then not only might the unrealized value be realized, but innovation will continue and again become self-sustaining. The ruggedness of the technological landscape captures the unevenness of innovation, exposing the conditions when innovation can stop. But in so doing, it also exposes the opportunities for it to be restarted.

In this section we develop and explore a policy intervention to correct this market failure. We consider a slight variant on our model. Specifically, we add a fixed component to R&D costs and ask how a policy that subsidizes this fixed cost affects innovation and welfare. We suppose that the government has full knowledge of the state of the market and offers the intervention *only* when innovation would otherwise stop. As firms still must pay the marginal cost of R&D, it is profitable for the firm to accept the intervention only if the frontier product's quality is positive and the new product creates positive value.[21]

Table 2 reports the results for the subsidy at different levels of the fixed cost of R&D. The simulations use the benchmark parameters from the previous section. As intuition suggests, the fixed cost of R&D slows down innovation at the frontier, reducing the number of new products introduced and lowering welfare (compare to the first column of Table 1 which

[21] There are many ways to design an intervention program. This intervention is simple and plausible and offers two additional benefits. First, when either frontier or niche innovation occurs, the product chosen is the same regardless of the fixed cost. The difference is that the fixed cost makes frontier innovation less attractive, turning firms back to niche innovation earlier or stopping innovation altogether. Thus, the effect of the subsidy is to change the type of innovation without changing the underlying incentives of which innovation to choose given that type. A second advantage is that this approach avoids the incentive compatibility problem that would arise if firms received a lump sum to avoid variable R&D costs as they would then simply shirk.

Fixed Cost	20	40	80	200
New products without subsidy	45.1	41.4	35.6	0
New products with subsidy	47.4	45.4	44.5	42.4
Markets with intervention	10.4%	16.7%	26.5%	100%
Products subsidized conditional on intervention	1.24	1.35	1.49	1.79
Welfare without subsidy	10973	10362	9178	125
Welfare gain conditional on intervention	3238	3913	4402	8571
Subsidy cost conditional on intervention	8.4	17.7	37.8	153.6

Table 2: Policy Intervention: $v_0 = 5, \sigma = 1, \tau = 10, c = 0.1$, $\gamma = 0.001$, $\delta = 0.9$, 200,000 simulations.

represents the zero fixed cost case).

The policy intervention creates substantial societal welfare. For a fixed R&D cost of 20, the government intervenes in only 10.4% of markets and on average 1.24 times when it does. The average subsidy costs 8.4 units of welfare, which is dwarfed by the welfare gain that is orders of magnitude higher at 3238. This represents a return-on-investment of close to 40,000%. Across all markets (subsidized and unsubsidized), the subsidy represents approximately a 3% increase in expected welfare.

The benefits of the intervention increase from there for larger fixed costs of R&D. For higher fixed costs, the rate of government intervention is higher, both in the fraction of markets and the interventions per market. The welfare gains are higher too and off a lower base. The return on investment is lower, as the cost of intervention increases at a higher rate, though it is still attractively high.

Fixed costs of 200 provide a particularly illuminating case. Without government intervention, there is no innovation in the market. (The welfare without intervention of 125 comes exclusively from product l_0.) Government intervention essentially makes the market by inducing a large amount of innovation. Much of that innovation is self-sustained, as the government subsidizes a product on average only 1.79 times, with the remaining 40 odd innovations introduced because they are profitable for the firms to do so.

These results suggest that the situation in Figure 9 is not unusual. That on a rugged technological landscape it is likely that innovation will stop prematurely, and that a little government nudging can restart the process such that innovation is again self-sustaining.

7 Conclusion

The relationship between innovation and competition is fundamental to economic growth. Our paper is one further step in the development of a foundation for this relationship. We develop a model of innovation in which the technological landscape is rugged, with ups and downs and pitfalls that can snag the innovative process. In providing a microfoundation that contains a sense of distance and direction, we are able to provide insight not only into whether firms innovate, but the type, the direction, and the boldness of innovation when they do. We provide the conditions for when innovation stops, and show how a targeted policy intervention can restart a stalled innovation process and add considerable value to society.

The essential ingredient underlying our results is the rugged landscape. We employ the Brownian motion as it captures this ruggedness in a tractable and neutral way. Some markets may not fit the linearity of drift and the independent and identical increments that define the Brownian motion. The rich variety of stochastic processes beyond the Brownian motion can be fitted to particular contexts. For instance, the mean reverting property of the Ornstein-Uhlenbeck process may better capture technologies where breakthroughs (and failures) are more localized.

The model of innovation and competition that we build on top of this landscape is parsimonious and relatively frictionless. It is easy to imagine other frictions and forces that shape innovation and these can be incorporated into the model. An important pair of extensions is to allow for far-sighted firms and for non-experience goods. If we slow or remove the competitive fringe, innovation becomes more valuable the longer a firm's horizon is, and as we saw earlier, firms will seek out niche innovations that remain active for longer terms. Similarly, non-experience goods may induce greater experimentation as profits would be convex in realized quality. Failures can be abandoned immediately, whereas successes can persist. Intuition suggests that in both variations the positive role of government intervention into the market would be enhanced.

The model can also be extended at a more foundational level. Our model adopts a particular conception of innovation in which the set of available innovations is known, and that all that is unknown are the realized qualities. Kuhn (1962) famously distinguished between science and technology, in his terms the the creation of knowledge and the practical uses of it, respectively.[22] Our model captures the technology side of this dichotomy. To extend the model toward the science side, several possibilities present themselves. A first step is to suppose that the firms do not know the drift and variance of the technology space (or, indeed, the generating process). The deeper goal is to formulate a discovery process that identifies the products themselves. Our model would then provide the guide to how those discoveries make it into new product innovations and drive economic growth.

[22]This is analogous to Kuznets' (1962) notions of invention and discovery.

8 Proofs

Proof of Proposition 1: We proceed in three steps.

Step 1: The first step is to note that at the entrant's optimal location, the market between products l_0 and the entrant's product l_1 has to be covered, that is, there cannot be any consumers located between l_0 and l_1 who consume neither product when both are offered at a zero price. Suppose, to the contrary, that the market is not covered. The entrant can then generate the same sales at lower R&D cost by moving its location slightly to the left.

Step 2. The second step is to work out the entrant's profits if the market between products l_0 and l_1 is covered. There are three critical consumers. Consumer $(l_0 + l_1)/2$ is indifferent between the two products if they are both offered at zero price. Consumer $l_0 + \tau v(l_0)$ is the right-most consumer who is indifferent between consuming l_0 at zero price and consuming no product. And consumer $l_1 + \tau v(l_0)$ is the right-most consumer who is indifferent between consuming l_1 at zero price and consuming no product, where we used $\mathrm{E}[v(l_1)|l_0] = v(l_0)$. We can then write the first entrant's profits as

$$\begin{aligned}\pi_1(l_1) &= \int_{\frac{1}{2}(l_0+l_1)}^{l_1} \left(v(l_0) - \frac{1}{\tau}(l_1 - s)\right) \mathrm{d}s \\ &\quad + \int_{l_1}^{l_1+\tau v(l_0)} \left(v(l_0) - \frac{1}{\tau}(s - l_1)\right) \mathrm{d}s \\ &\quad - \int_{\frac{1}{2}(l_0+l_1)}^{l_0+\tau v(l_0)} \left(v(l_0) - \frac{1}{\tau}(s - l_0)\right) \mathrm{d}s - \frac{1}{2}c(l_1 - l_0)^2,\end{aligned}$$

where the first two terms are the gross value created by the entrant, the third is the value captured by consumers, and the last is the cost of developing the product. Evaluating this expression we have that the entrant's profits are given by

$$\pi_1(l_1) = \tau v(l_0)^2 - \frac{1}{4\tau}(2\tau v(l_0) - (l_1 - l_0))^2 - \frac{1}{2}c(l_1 - l_0)^2. \tag{6}$$

Step 3: The third step is to maximize (6) with respect to location l_1, which delivers the optimal location

$$l_1^* = l_0 + 2\tau \frac{v(l_0)}{1 + 2c\tau}.$$

Substituting back into (6) shows that at the optimal location the entrant's profits are

$$\pi_1(l_1^*) = \tau v(l_0)^2 - \frac{2c\tau^2}{1 + 2c\tau} v(l_0)^2 > 0. \quad \blacksquare$$

Proof of Proposition 2: It is immediate that if $v\left(l_t^f\right) \leq 0$, the entrant is better off locating at l_t^f than strictly to its right. For the rest of this proof suppose, therefore, that $v\left(l_t^f\right) > 0$. We proceed in four steps.

Step 1: Analogous to the first step in Proposition 1, at the entrant's optimal location the market between l_t^a and l_t has to be covered. If it were not, the entrant could generate the same sales at lower R&D costs by moving slightly to the left.

Step 2: The entrant will not have any customers if it locates so close to the frontier that even its ideal customer l_t prefers the right-most active product l_t^a to the entrant's product l_t when both are offered at zero price, that is, if

$$v\left(l_t^a\right) - \frac{1}{\tau}\left(l_t - l_t^a\right) > \mathrm{E}\left[v\left(l_t\right)\middle| l_t^f\right],$$

or, equivalently,

$$l_t < l_t^a + \tau\left(v\left(l_t^a\right) - v\left(l_t^f\right)\right), \tag{7}$$

where we used $\mathrm{E}\left[v\left(l_t\right)\middle| l_t^f\right] = v\left(l_t^f\right)$.

Step 3: The next step is to work out the entrant's profits for locations l_t such that the market between l_t^a and l_t is covered and there are some customers who buy the entrant's product, that is, (7) is not satisfied. There then exists a consumer located between l_t^a and l_t who is indifferent between the two products when they are both offered at zero price. The location x_t of this consumer satisfies

$$v\left(l_t^a\right) - \frac{1}{\tau}\left(x_t - l_t^a\right) = v\left(l_t^f\right) - \frac{1}{\tau}\left(l_t - x_t\right)$$

and is thus given by

$$x_t = \frac{1}{2}\left(l_t^a + l_t\right) + \frac{1}{2}\tau\left(v\left(l_t^a\right) - v\left(l_t^f\right)\right).$$

This consumer is the left-most consumer who consumes the entrant's product. The right-most consumer who does so is given by $l_t + \tau v\left(l_t^f\right)$. From this, the entrant's profits are

$$\begin{aligned}\pi_t\left(l_t\right) \;=\; & \int_{x_t}^{l_t}\left(v\left(l_t^f\right) - \frac{1}{\tau}\left(l_t - s\right)\right)\mathrm{d}s + \int_{l_t}^{l_t+\tau v\left(l_t^f\right)}\left(v\left(l_t^f\right) - \frac{1}{\tau}\left(s - l_t\right)\right)\mathrm{d}s \\ & - \int_{x_t}^{l_t^a+\tau v\left(l_t^a\right)}\left(v\left(l_t^a\right) - \frac{1}{\tau}\left(s - l_t^a\right)\right)\mathrm{d}s - \frac{1}{2}c\left(l_t - l_t^f\right)^2,\end{aligned}$$

where the first two terms are the value created by the entrant, the third term is the value captured by the entrant's customers, and the last term are the costs of developing the entrant's product. Simplifying the above expression, we have

$$\pi_t(l_t) = \tau v\left(l_t^f\right)^2 - \frac{\tau}{4}\left(v(l_t^a) + v\left(l_t^f\right) - \frac{1}{\tau}(l_t - l_t^a)\right)^2 - \frac{1}{2}c\left(l_t - l_t^f\right)^2. \tag{8}$$

Step 4: The location that maximizes (8) is given by

$$\widehat{l_t} = l_t^f + \frac{\tau}{1+2c\tau}\left(v\left(l_t^f\right) + v(l_t^a) - \frac{1}{\tau}\left(l_t^f - l_t^a\right)\right).$$

It is routine to confirm that the market between l_t^a and $\widehat{l_t}$ is covered. For $\widehat{l_t}$ to be the optimal location to the right of l_t, it is not enough that the entrant has customers at this location, that is, that $\widehat{l_t} \geq l_t^a + \tau\left(v(l_t^a) - v\left(l_t^f\right)\right)$. Instead, its revenue has to be enough to at least cover the development costs. Substituting $\widehat{l_t}$ into (8) we have

$$\pi_t\left(\widehat{l_t}\right) = \tau v\left(l_t^f\right)^2 - \frac{1}{2}\frac{c\tau^2}{1+2c\tau}\left(u\left(l_t^f, l_t^a\right) + v\left(l_t^f\right)\right)^2.$$

This expression is increasing in $v\left(l_t^f\right)$ and equal to zero for $v\left(l_t^f\right) = \kappa \cdot u\left(l_t^f, l_t^a\right)$, where

$$\kappa \equiv \frac{\sqrt{\frac{c\tau}{2(1+2c\tau)}}}{1 - \sqrt{\frac{c\tau}{2(1+2c\tau)}}}. \tag{9}$$

In summary, if $v\left(l_t^f\right) > \kappa \cdot u\left(l_t^f, l_t^a\right)$, frontier innovation is profitable and the optimal location to the right of l_t is given by $l_t^{f*} = \widehat{l_t}$. If, instead, $v_f \leq \kappa \cdot u\left(l_t^f, l_t^a\right)$, frontier innovation is not profitable. ■

Proof of Proposition 3: In period t the entrant locates at some some $l_t \in [a,b]$. The left-most consumer $\underline{s}_l$ who buys the new product satisfies

$$\mathrm{E}(|v(a)\,\mathcal{E}_t) - \frac{1}{\tau}\left(\underline{s}_{l_t} - a\right) = \mathrm{E}(|v(l_t)\,\mathcal{E}_t) - \frac{1}{\tau}\left(l_t - \underline{s}_{l_t}\right)$$

and is thus given by

$$\underline{s}_{l_t} = a + \frac{1}{2}(l_t - a)(1 - \tau\beta(a,b)), \tag{10}$$

where

$$\beta(a,b) \equiv \frac{\mathrm{E}(|v(b)\,\mathcal{E}_t) - \mathrm{E}(|v(a)\,\mathcal{E}_t)}{b-a}.$$

Similarly, the right-most consumer who buys the new product is given by

$$\bar{s}_{l_t} = b - \frac{1}{2}(b - l_t)(1 + \tau\beta(a,b)). \tag{11}$$

Finally, the consumer s_m who is indifferent between the neighboring active products is given by

$$\mathrm{E}(|v(a)\ \mathcal{E}_t) - \frac{1}{\tau}(s_m - a) = \mathrm{E}(|v(b)\ \mathcal{E}_t) - \frac{1}{\tau}(b - s_m)$$

and is thus given by

$$s_{ab} = a + \frac{1}{2}(b - a)(1 - \tau\beta(a,b)) \tag{12}$$

We can then write the entrant's profits as

$$\begin{aligned}
\pi_t(l_t) &= \int_{\underline{s}_{l_t}}^{l_t} \mathrm{E}(|v(l_t)\ \mathcal{E}_t) - \frac{1}{\tau}(l_t - s)\,\mathrm{d}s + \int_{l_t}^{\bar{s}_{l_t}} \mathrm{E}(|v(l_t)\ \mathcal{E}_t) - \frac{1}{\tau}(s - l_t)\,\mathrm{d}s \\
&\quad - \int_{\underline{s}_{l_t}}^{s_{ab}} \mathrm{E}(|v(a)\ \mathcal{E}_t) - \frac{1}{\tau}(s - a)\,\mathrm{d}s - \int_{s_{ab}}^{\bar{s}_{l_t}} \mathrm{E}(|v(b)\ \mathcal{E}_t) - \frac{1}{\tau}(b - s)\,\mathrm{d}s \\
&= \frac{1}{2\tau}(l_t - a)(b - l_t)\left(1 - (\tau\beta(a,b))^2\right).
\end{aligned}$$

The optimal location is then given by $l_t^{n*}(a,b) = (a+b)/2$ and optimal profits are given by

$$\pi_t\left(l_t^{n*}(a,b)\right) = \frac{1}{8\tau}(b-a)^2\left(1 - (\tau\beta(a,b))^2\right). \quad \blacksquare$$

The following lemma is useful for Proposition 4

Lemma 1 *Suppose frontier innovation is optimal for firm $t \geq 1$. If $c \geq \frac{2}{3\tau}$, it cannot be optimal for firm $t+1$ to locate in $\left[l_t^f, l_t^*\right]$.*

Proof: It is immediate that firm $t+1$ will not enter at l_t^f or l_t^*, since competition with other firms at the same location would drive its profits down to zero. The remainder of the proof shows that firm $t+1$ will not enter in $\left(l_t^f, l_t^*\right)$ either. We proceed in three steps.

Step 1: We know from Propositions 1 and 2 that since it is optimal for firm t to engage in frontier innovation, it locates at $l_t^* = {l_t}^f + \frac{\tau}{1+2c\tau}\left(v\left(l_t^f\right) + u\left(l_t^f, l_t^a\right)\right)$. In period $t+1$, the frontier product is therefore given by

$$l_{t+1}^f = l_t^* = {l_t}^f + \frac{\tau}{1+2c\tau}\left(v\left(l_t^f\right) + u\left(l_t^f, l_t^a\right)\right). \tag{13}$$

It is immediate that it cannot be optimal for firm $t+1$ to locate in $\left(l_t^f, l_t^*\right)$ if the frontier product l_{t+1}^f is not active. For the remainder of the proof, we, therefore, assume that l_{t+1}^f is active, that is, $v\left(l_{t+1}^f\right) \geq u\left(l_t^f, l_t^a\right) - \frac{1}{\tau}\left(l_{t+1}^f - l_t^f\right)$. Using (13) to substitute for $l_{t+1}^f - l_t^f$ we can rewrite this condition as

$$v\left(l_{t+1}^f\right) \geq \frac{1}{1+2c\tau}\left(2c\tau u\left(l_t^f, l_t^a\right) - v\left(l_t^f\right)\right). \tag{14}$$

It is also immediate that it cannot be optimal for firm $t+1$ to locate in $\left(l_t^f, l_t^*\right)$ if the quality of l_{t+1}^f is so high that even consumer l_t^f gets a higher gross quality from l_{t+1}^f than from l_t^f. For the remainder of the proof we, therefore, assume that $v\left(l_t^f\right) \geq v\left(l_{t+1}^f\right) - \frac{1}{\tau}\left(l_{t+1}^f - l_t^f\right)$ or, equivalently,

$$\frac{1}{\tau} \geq \frac{v\left(l_{t+1}^f\right) - v\left(l_t^f\right)}{\left(l_{t+1}^f - l_t^f\right)} = \beta\left(l_t^f, l_{t+1}^f\right). \tag{15}$$

We will be using both (14) and (15) below.

Step 2: The second step in the proof is to bound the gross utility $u\left(l_t^f, l_t^a\right) = v\left(l_t^a\right) - \frac{1}{\tau}\left(l_t^f - l_t^a\right)$ that consumer l_t^f gets from consuming the right-most active product l_t^a. We will be using these bounds in Step 3 below.

If the frontier product l_t^f is active, it is the right-most active product, and thus $u\left(l_t^f, l_t^a\right) = u\left(l_t^f, l_t^f\right) = v\left(l_t^f\right)$. If, instead, the frontier product is not active, we have $l_t^a < l_t^f$. Consumer l_t^f then realizes a higher gross utility from consuming l_t^a than from consuming l_t^f. The lower bound of $u\left(l_t^f, l_t^a\right)$ is, therefore, given by $v\left(l_t^f\right)$, that is, $u\left(l_t^f, l_t^a\right) \geq v\left(l_t^f\right)$.

To obtain an upper bound, we use the fact that frontier innovation is optimal for firm t. It then follows from Proposition 2, and the definition of κ in (9), that $v\left(l_t^f\right) \geq \kappa u\left(l_t^f, l_t^a\right)$. Rearranging this expression, we have that

$$u\left(l_t^f, l_t^a\right) \leq \left(\sqrt{2\frac{1+2c\tau}{c\tau}} - 1\right) v\left(l_t^f\right) \equiv \overline{u}_t. \tag{16}$$

If consumer l_t^f's gross utility from consuming l_t^a were higher than this upper bound $\overline{u}_t$, it would not be optimal for firm t to engage in frontier innovation, which contradicts the assumption in the proposition.

Step 3: We now turn to the main part of the proof, which compares profits under niche and frontier innovation. We know from Propositions 1 and 2 that if firm $t+1$ engages in frontier innovation, its profits are given by

$$\pi_{t+1}\left(l_{t+1}^f\right) = \frac{\tau}{1+2c\tau} v\left(l_{t+1}^f\right)^2. \tag{17}$$

Moreover, we know from Proposition 3 that if firm $t+1$ engages in niche innovation, its profits

are $\pi_t\left(l_t^{n*}\left(a,l_{t+1}^f\right)\right)=\frac{1}{8\tau}\left(l_{t+1}^f-a\right)^2\left(1-\left(\tau\beta\left(a,l_{t+1}^f\right)\right)^2\right)$, where
$a=l_t^f+\frac{\tau}{1+\tau\beta\left(a,l_{t+1}^f\right)}\left(v\left(l_t^a\right)-v\left(l_t^f\right)\right)$ is the left-most point in the viable niche $\left[a,l_{t+1}^f\right]$. Substituting this expression for a into niche profits, we have

$$\pi_t\left(l_t^{n*}\left(a,l_{t+1}^f\right)\right)=\frac{\tau}{8}\left(v\left(l_{t+1}^f\right)-\frac{1}{1+2c\tau}\left(2c\tau u\left(l_t^f,l_t^a\right)-v\left(l_t^f\right)\right)\right)^2\frac{1-\tau\beta\left(a,l_{t+1}^f\right)}{1+\tau\beta\left(a,l_{t+1}^f\right)}$$

Differentiating with respect to $u\left(l_t^f,l_t^a\right)$ we have

$$\begin{aligned}\frac{\mathrm{d}\pi_t\left(l_t^{n*}\left(a,l_{t+1}^f\right)\right)}{\mathrm{d}u\left(l_t^f,l_t^a\right)} &= -\frac{c\tau}{2c\tau+1}\frac{\tau}{2}\frac{1-\tau\beta\left(a,l_{t+1}^f\right)}{1+\tau\beta\left(a,l_{t+1}^f\right)}\left(v\left(l_{t+1}^f\right)-\frac{1}{1+2c\tau}\left(2c\tau u\left(l_t^f,l_t^a\right)-v\left(l_t^f\right)\right)\right)\\ &\quad+\frac{\tau}{8}\left(v\left(l_{t+1}^f\right)-\frac{1}{1+2c\tau}\left(2c\tau u\left(l_t^f,l_t^a\right)-v\left(l_t^f\right)\right)\right)^2\\ &\quad\times\frac{2\tau}{\left(1+\tau\beta\left(a,l_{t+1}^f\right)\right)^2}\frac{\beta\left(a,l_{t+1}^f\right)}{\left(u\left(l_t^f,l_t^a\right)+v\left(l_t^f\right)\right)^2}.\end{aligned}$$

The first term on the right-hand side is negative because product l_{t+1}^f is active, that is, condition (14) is satisfied, and because (15) is assumed to hold. The sign of the second term is the same as the sign of the slope $\beta\left(a,l_{t+1}^f\right)$.

Step 3a: Suppose first that $\beta\left(a,l_{t+1}^f\right)\leq 0$, in which case $\pi_t\left(l_t^{n*}\left(a,l_{t+1}^f\right)\right)$ is decreasing in $u\left(l_t^f,l_t^a\right)$. We know from Step 1 that the lowest value $u\left(l_t^f,l_t^a\right)$ can take is $v\left(l_t^f\right)$. We then have

$$\begin{aligned}\pi_t\left(l_t^{n*}\left(a,l_{t+1}^f\right)\right) &= \frac{\tau}{8}\left(\frac{1}{1+2c\tau}\left(v\left(l_t^f\right)+u\left(l_t^f,l_t^a\right)\right)-\frac{1}{1+\tau\beta}\left(u\left(l_t^f,l_t^a\right)-v\left(l_t^f\right)\right)\right)^2 \qquad (18)\\ &\quad\times\left(1-\left(\tau\beta\left(a,l_{t+1}^f\right)\right)^2\right)\\ &\leq \frac{\tau}{2\left(1+2c\tau\right)^2}\left(1-\left(\tau\beta\left(a,l_{t+1}^f\right)\right)^2\right)v\left(l_t^f\right)^2,\end{aligned}$$

where the second line follows from the first by substituting $v\left(l_t^f\right)$ for $u\left(l_t^f,l_t^a\right)$. Using the expression for frontier profits (17) we then have that frontier innovation dominates niche innovation if

$$\frac{\tau}{1+2c\tau}v\left(l_{t+1}^{f}\right)^{2}\geq\frac{\tau}{2\left(1+2c\tau\right)}\left(1-\left(\tau\beta\left(a,l_{t+1}^{f}\right)\right)^{2}\right)v\left(l_{t}^{f}\right)^{2}.$$

Substituting for

$$v\left(l_{t+1}^{f}\right)=v\left(l_{t}^{f}\right)+\beta\left(a,l_{t+1}^{f}\right)\frac{2\tau}{2c\tau+1}v\left(l_{t}^{f}\right)$$

this becomes

$$2\left(1+2c\tau+2\tau\beta\left(a,l_{t+1}^{f}\right)\right)^{2}-\left(1+2c\tau\right)\left(1-\left(\tau\beta\left(a,l_{t+1}^{f}\right)\right)^{2}\right)\geq 0. \tag{19}$$

The left-hand side is convex in $\beta\left(a,l_{t+1}^{f}\right)$ and takes its smallest value for $\beta\left(a,l_{t+1}^{f}\right)=-\frac{8c\tau+4}{9\tau+2c\tau^{2}}$. Evaluated for this value of $\beta\left(a,l_{t+1}^{f}\right)$, the left-hand side of (19) is positive for all $c\geq\frac{2}{3\tau}$, which is the constraint on c assumed in the lemma.

Step 3b: Suppose next that $\beta\left(a,l_{t+1}^{f}\right)\geq 0$, that is, $v\left(l_{t+1}^{f}\right)\geq v\left(l_{t}^{f}\right)$. Recall that if firm $t+1$ engages in frontier innovation, its profits are given by (17). Since $v\left(l_{t+1}^{f}\right)\geq v\left(l_{t}^{f}\right)$, this is bounded from below by $\underline{\pi}_{t+1}\equiv\frac{\tau}{1+2c\tau}v\left(l_{t}^{f}\right)^{2}$.

Next, recall that $\overline{u}_{t}$, the upper bound of $u\left(l_{t}^{f},l_{t}^{a}\right)$, is defined in (14). We can then rewrite the profits (18) that firm $t+1$ makes from niche innovation as

$$\begin{aligned}
\pi_{t}\left(l_{t}^{n*}\left(a,l_{t+1}^{f}\right)\right) &= \frac{1}{8\tau}\left(\begin{array}{c}\frac{\tau}{1+2c\tau}\left(v\left(l_{t}^{f}\right)+\overline{u}_{t}\right)\\ -\left(\frac{\tau}{1+2c\tau}\left(\overline{u}_{t}-u\left(l_{t}^{f},l_{t}^{a}\right)\right)+\frac{\tau}{1+\tau\beta}\left(u\left(l_{t}^{f},l_{t}^{a}\right)-v\left(l_{t}^{f}\right)\right)\right)\end{array}\right)^{2}\\
&\quad\times\left(1-\left(\tau\beta\left(a,l_{t+1}^{f}\right)\right)^{2}\right)\\
&\leq \frac{1}{8\tau}\left(\frac{\tau}{1+2c\tau}\left(v\left(l_{t}^{f}\right)+\overline{u}_{t}\right)\right)^{2}\left(1-\left(\tau\beta\left(a,l_{t+1}^{f}\right)\right)^{2}\right)\\
&\leq \frac{1}{8\tau}\left(\frac{\tau}{1+2c\tau}\left(v\left(l_{t}^{f}\right)+\overline{u}_{t}\right)\right)^{2}\equiv\overline{\pi}_{t+1}^{n},
\end{aligned}$$

where the first inequality follows from the facts, established in Step 2, that $u\left(l_{t}^{f},l_{t}^{a}\right)$ is bounded above by $\overline{u}_{t}$ and below by $v\left(l_{t}^{f}\right)$ and the second inequality follows from (15). Subtracting the upper bound on niche profits $\overline{\pi}_{t+1}^{n}$ from the lower bound on frontier profits $\underline{\pi}_{t+1}^{f}$ we get $\underline{\pi}_{t+1}^{f}-\overline{\pi}_{t+1}^{n}=\frac{1}{4}\frac{4c\tau-1}{c(2c\tau+1)}v\left(l_{t}^{f}\right)^{2}>0$, which completes the proof. ■

Proof of Proposition 4: Consider first period $t=2$. We know from Lemma 1 that firm t will not locate at any location in $[l_0,l_1]$. From Proposition 2, firm 2 then engages in frontier innovation if $v\left(l_2^f\right) > \kappa u\left(l_2^f, l_2^a\right)$, and it will not enter otherwise. Moreover, if firm 2 does not enter, neither will any future entrant. In period 2, therefore, $\overline{v}_2^f = \kappa u\left(l_2^f, l_2^a\right)$.

Next, consider any period $t > 2$. We know from Lemma A1 that in any period t the two right-most products are given by l_{t-1}^f and l_t^f and that firm t will not locate in $\left[l_{t-1}^f, l_t^f\right]$. Let $\Pi_t \geq 0$ denote the highest profit firm t can realize by locating in $\left[l_0^f, l_{t-1}^f\right)$. From Proposition 2, frontier innovation is profitable if $v\left(l_t^f\right) > \kappa u\left(l_t^f, l_t^a\right)$. Moreover, when it is profitable, it generates profit

$$\pi_t\left(l_t^{f*}\right) = \tau v\left(l_t^f\right)^2 - \frac{1}{2}\frac{c\tau^2}{1+2c\tau}\left(u\left(l_t^f, l_t^a\right) + v\left(l_t^f\right)\right)^2,$$

which is increasing in $v\left(l_t^f\right)$. We can then define $\overline{v}_t^f = \kappa u\left(l_t^f, l_t^a\right)$ if $\Pi_t = 0$ and $\overline{v}_t^f > \kappa u\left(l_t^f, l_t^a\right)$ as the value of $v\left(l_t^f\right)$ for which $\pi_t\left(l_t^{f*}\right) = \Pi_t$.

Finally, notice that if frontier innovation is not profitable in period t, it will also not be profitable for any firm $t' > t$. If, however, in period t, frontier innovation is profitable, but not as profitable as niche innovation, frontier innovation innovation may again be optimal for some firm $t' > t$. The reason is that niche innovation by firm t, and possibly other firms, may bring down the profits $\Pi_{t'}$ that firm t' can generate by engaging in niche innovation. ■

Proof of Corollary 1: Innovation to the left of l_t never occurs if all incumbent products are in the competitive shadow of l_t as then every niche is non-viable. A necessary condition is that $v(l_t)$ is higher than all incumbent products. As l_t's shadow increases without bound in quality, the threshold must satisfy a cut-point, and the result follows. ■

Proof of Proposition 5: Let l_A denote the closest active product to the left of a and l_B denote the closest active product to the right of b. Next, define

$$\underline{v}_t^n(a,b) = \min\left[v(l_A) - \frac{1}{\tau}\left(\frac{1}{2}(a+b) - l_A\right), v(l_B) - \frac{1}{\tau}\left(l_B - \frac{1}{2}(a+b)\right)\right].$$

Note that $l_t^* = \frac{1}{2}(a+b)$ and suppose $v(l_t^*) \leq \underline{v}_t^n(a,b)$. In period $t+1$, the expected gross utility that any consumer $s \in [a,b]$ realizes from consuming any product $l \in [a,b]$ is either smaller than the gross utility the consumer realizes from consuming l_A or it is smaller than the gross utility the consumer realizes from consuming l_B. As such, there exists no viable niche in $[a,b]$ in period $t+1$ or in any future period.

Next, let l_a denote the closest existing product to the left of a and l_b denote the closest existing product to the right of b. Define

$$\overline{v}_t^n(a,b) = \max\left[v(l_a) + \frac{1}{\tau}\left(\frac{1}{2}(a+b) - l_a\right), v(l_b) + \frac{1}{\tau}\left(l_b - \frac{1}{2}(a+b)\right)\right].$$

Suppose that $v(l_t^*) \geq \overline{v}_t^n(a,b)$. In period $t+1$, the expected gross utility that any consumer $s \in [a,b]$ realizes from consuming any product $l \in [a,b]$ is smaller than the gross utility the consumer realizes from consuming $l_t^* = \frac{1}{2}(a+b)$. As such, there exists no viable niche in $[a,b]$ in period $t+1$ or in any future period.

Finally, suppose $\underline{v}_t^n(a,b) < v(l_t^*) < \overline{v}_t^n(a,b)$. In period $t+1$, there then exists at least one type of consumer $s \in [a,b]$ for whom the expected gross utility from consuming product $l = s$ is strictly higher than the gross utility from consuming either l_a or l_b. As such, there exists a viable niche in $[a,b]$ in period $t+1$. ■

Proof of Corollary 2: The result follows if $v(l_t)$ relative to the frontier product satisfies case (i) of Proposition 2 and the requirements of Corollary 1 are satisfied. ■

Proof of Proposition 6: Consider the following hypothetical. Take the interval $[0,l]$ and in each period slice one of the longest niches into two equal niches such that after t periods there are $t+1$ niches, and, for each positive integer q, after $2^q - 1$ periods the 2^q niches are of equal length $\frac{l}{2^q}$. For each innovation, consider a niche dead if the realized outcome in either direction further from the mean than $\frac{1}{\tau}$ times the niche's length. On the Brownian path, the realized outcome on a niche of length x is normally distributed with variance $\frac{x}{4}$. Thus, the probability the niche dies is no longer viable from one period to the next is bounded below by:

$$CDF\left(-\frac{x}{\tau}\right) + \left(1 - CDF\left(\frac{x}{\tau}\right)\right) = 1 - erf\left(\frac{\sqrt{2x}}{\tau}\right), \tag{20}$$

which is strictly decreasing in x and approaches 1 as $x \to 0^+$. After some number of periods, t', this probability for each surviving niche is at least $1-\epsilon$ for some ϵ small. For $t > t'$, the expected number of surviving niches decreases each period and approaches 0.

The proposition follows as the probability that innovation on the niche $[0,l]$ continues in equilibrium is bounded above by the hypothetical process just described. Proposition 3 implies that innovation slices a niche into equal lengths, as in the hypothetical. The exception is when a niche is not entirely optimal. Thus, in period t, the length of a niche on which innovation occurs is weakly shortly than in the hypothetical, and (20) is decreasing in x. Finally, Proposition 5 implies that a realized outcome $\frac{x}{4}$ or further from the mean is sufficient to stop innovation in the niche, regardless of the slope of the niche or whether it is entirely viable (and a high outcome may stop innovation in other niches). ■

References

Acemoglu, Daron. 1998. "Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality." *Quarterly Journal of Economics* 113(4):1055–1089.

Aghion, Philippe, Christopher Harris, Peter Howitt and John Vickers. 2001. "Competition, Imitation and Growth with Step-by-Step Innovation." *Review of Economic Studies* 68(3):467–492.

Aghion, Philippe, Nick Bloom, Richard Blundell, Rachel Griffith and Peter Howitt. 2005. "Competition and Innovation: An Inverted-U Relationship." *Quarterly Journal of Economics* 120(2):701–728.

Aghion, Philippe and Peter Howitt. 1992. "A Model of Growth Through Creative Destruction." *Econometrica* 60(2):323–351.

Arcidiacono, Peter, Patrick Bayer, Jason R. Blevins and Paul B. Ellickson. 2016. "Estimation of Dynamic Discrete Choice Models in Continuous Time with an Application to Retail Competition." *Review of Economic Studies* 83(3):889–931.

Braguinsky, Serguey, Atsushi Ohyama, Tetsuji Okazaki and Chad Syverson. 2021. "Product Innovation, Product Diversification, and Firm Growth: Evidence from Japan's Early Industrialization." *American Economic Review* 111(12):3795–3826.

Bryan, Kevin A. and Jorge Lemus. 2017. "The direction of innovation." *Journal of Economic Theory* 172:247–272.

Cabral, Luis M. B. 2003. "R&D Competition When Firms Choose Variance." *Journal of Economics and Management Strategy* 12(1):139–150.

Callander, Steven. 2011. "Seaching and Learning by Trial and Error." *American Economic Review* 101(6):2277–2308.

Callander, Steven and Niko Matouschek. 2019. "The Risk of Failure: Trial and Error Learning and Long-Run Performance." *American Economic Journal: Microeconomics* 11(1):44–78.

Callander, Steven and Niko Matouschek. 2022. "The Novelty of Innovation: Competition, Disruption, and Antitrust Policy." *Management Science* 68(1):37–51.

Chamberlin, Edward H. 1933. *The Theory of Monopolistic Competition*. Harvard University Press.

DiMasi, Jospeh A. and Laura B. Faden. 2011. "Competitiveness in Follow-on Drug R&D: a Race or Imitation?" *Nature Reviews* 10:23–27.

Dixit, Avinash K. and Joseph E. Stiglitz. 1977. "Monopolistic Competition and Optimum Product Diversity." *American Economic Review* 67(3):297–308.

Garcia-Macia, Daniel, Chang-Tai Hsieh and Peter J. Klenow. 2019. "How Destructive is Innovation?" *Econometrica* 87(5):1507–1541.

Garfagnini, Umberto and Bruno Strulovici. 2016. "Social Experimentation with Interdependent and Expanding Technologies." *Review of Economic Studies* 83(4):1579–1613.

Granja, Joao and Sara Moreira. 2021. "Product Innovation and Credit Market Disruptions." *Review of Financial Studies* forthcoming.

Hannan, Michael T. and John Freeman. 1977. "The Population Ecology of Organizations." *The American Journal of Sociology* 82(5):929–964.

Harris, Christopher and John Vickers. 1987. "Racing with Uncertainty." *Review of Economic Studies* 54(1):1–21.

Hotelling, Harold. 1929. "Stability in Competition." *Economic Journal* 39:41–57.

Kim, W. Chan and Renee Mauborgne. 2004. *Blue Ocean Strategy*. Harvard Business School Press.

Krieger, Joshua, Danielle Li and Dimitris Papanikolaou. 2021. "Missing Novelty in Drug Development." *The Review of Financial Studies* 35(2):636–679.
URL: *https://doi.org/10.1093/rfs/hhab024*

Kuhn, Thomas. 1962. Comment. In *The Rate and Direction of Inventive Activity: Economic and Social Factors*. Princeton, NJ: Princeton University Press pp. 379–384.

Kuznets, Simon. 1962. Inventive Activity: Problems of Definition and Measurement. In *The Rate and Direction of Inventive Activity: Economic and Social Factors*. Princeton, NJ: Princeton University Press pp. 19–43.

Lancaster, Kelvin. 1979. *Variety, Equity, and Efficiency*. Columbia University Press.

Lancaster, Kelvin. 1990. "The Economics of Product Variety: A Survey." *Marketing Science* 9(3):189–206.

Letina, Igor. 2016. "The road not taken: competition and the R& D portfolio." *The RAND Journal of Economics* 47:433–460.

Levinthal, Daniel A. 1997. "Adaption on Rugged Landscapes." *Management Science* 43(7):934–950.

Meagher, Kieron J. and Klaus G. Zauner. 2004. "Product Differentiation and Location Decisions Under Demand Uncertainty." *Journal of Economic Theory* 117:201–216.

Prescott, Edward C. and Michael Visscher. 1977. "Sequentil Location among Firms with Foresight." *The Bell Journal of Economics* 8(2):378–393.

Reinganum, Jennifer F. 1981. "Dynamic Games of Innovation." *Journal of Economic Theory* 25(1):21–41.

Reinganum, Jennifer F. 1982. "A Dynamic Game of R and D: Patent Protection and Competitive Behavior." *Econometrica* 50(3):671–688.

Reinganum, Jennifer F. 1983. "Unvertain Innovation and the Persistence of Monopoly." *The American Economic Review* 73(4):741–748.

Reinganum, Jennifer F. 1985. "Innovation and Industry Evolution." *Quarterly Journal of Economics* 100(1):81–99.

Rivkin, Jan W. 2000. "Imitation of Complex Strategies." *Management Science* 46(6):822–844.

Rogers, Everett. 1962. *Diffusion of Innovations*. Free Press of Glencoe.

Romer, Paul. 1990. "Endogenous Technological Change." *Journal of Political Economy* 98(5, Part 2):S71–S102.

Salop, Steven C. 1979. "Monopolistic Competition with Outside Goods." *The Bell Journal of Economics* 10(1):141–156.

Schumpeter, Joseph A. 1942. *Capitalism, Socialism, and Democracy*. New York: Harper and Brothers.

Segal, Ilya and Michael D. Whinston. 2007. "Antitrust in Innovative Industries." *American Economic Review* 97(5):1703–1730.

Wollmann, Thomas G. 2018. "Trucks Without Bailouts: Equilibrium Product Characteristics for Commmercial Vehicles." *American Economic Review* 108(6):1364–1406.