Innovation and Competition on a Rugged Technological Landscape*

Steven Callander† Nicolas S. Lambert‡ Niko Matouschek§

September 30, 2025

Abstract

We propose a model of dynamic spatial competition over a rugged technological landscape in which the quality of a product is learned only after it is introduced to the market. Firms enter the market sequentially, choosing not only whether to innovate but how to innovate. They can innovate beyond the frontier of existing products, outside the scope of the market, or they can innovate in a niche, nestling between incumbent products. The uncertainty about a new product's quality depends on the type of innovation a firm chooses 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 entry and product differentiation, narrower markets, and more intense competition than in a world of certainty.

JEL Codes: D21, L11.

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 (Schumpeter, 1942; Romer, 1990; Wollmann, 2018). The novelty of new products varies widely, however. Some are barely

*For helpful comments we thank Daron Acemoglu, Serguey Braguinsky, Christoph Carnehl, Ben Jones, Sarit Markovich, Sara Moreira, Spencer Pantoja, Bruno Strulovici, Takuo Sugaya, Can Urgun, and Leeat Yariv. The helpful feedback of the referees and Navin Kartik (the editor) improved the paper considerably.
†Stanford University, Stanford Graduate School of Business; sjc@stanford.edu.
‡University of Southern California, Department of Economics; lambertn@usc.edu.
§Northwestern University, Kellogg School of Management; n-matouschek@kellogg.northwestern.edu

tweaks on existing products, whereas others represent breakthrough innovations that depart radically from what has come before.[1]

For a firm entering a new market, therefore, the strategic problem 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 a firm resolves this problem has 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).

We introduce a framework that captures the strategic problem firms face and we explore its implications for innovation, competition, and industry dynamics. Formally, we meld Hotelling's (1929) classic model of spatial competition with a recent approach to modeling innovation. The set of potential products is the positive real line. Each product is horizontally differentiated from its neighbors a la Hotelling and vertically differentiated in quality. The quality of a product is determined by a rugged technological landscape that spans the product space and which we represent by the realized path of a Brownian motion (Callander, 2011). The technological landscape is unknown to the firms, with the quality of a product observed only after it is introduced to the market, and uncertainty about quality increasing in the novelty of the product itself. The framework captures and ties together the technological and competitive impact of innovation, consistent with the empirical evidence of Bloom, Schankerman and Van Reenen (2013).

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 and introduce a new product. A firm can innovate on the frontier outside the set of existing products (a blue ocean strategy), or it can innovate in a niche between existing products (a red ocean strategy).[2] If a firm enters the market, 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.

Our first set of results characterize optimal innovation by the firms. We show that the ruggedness of the technological landscape drives the innovation choices by firms. If the technological landscape were flat, an entering firm would face incumbent products of equal quality and the firm's only incentive would be to reduce competition. Following Hotelling's classic

[1] See Krieger, Li and Papanikolaou (2022) and Granja and Moreira (2023) for evidence from pharmaceuticals and insurance, respectively.

[2] 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." The outcome is ex ante uncertain and it need not be a success.

logic, the firm would differentiate itself horizontally by locating beyond the frontier of the existing market—a blue ocean strategy—and, on an unbounded product space, entry would continue indefinitely.

Innovation is very different on a rugged technological landscape. Even if uncertainty does not directly affect the preferences of the firms, the ruggedness of the technological landscape means that high quality products are more likely to lie in some parts of the product space than in other parts. This causes firms to turn away from the frontier when it is less promising and innovate in a niche instead. On a rugged technological landscape some firms adopt a red ocean strategy, trading off more intense competition for better prospects of a high quality product. On a rugged landscape, there is less product differentiation, reversals in both the direction and type of innovation, and the possibility that market entry eventually ends.

Niche and frontier innovation also differ in what drives profitability. At the frontier, the higher the quality of the boundary product, the more profitable it is for a new firm to innovate and the more novel is the product it chooses. In a niche, in contrast, profitability depends on the relative quality of the neighboring products more than the absolute quality. In a niche, innovation is less profitable when one of the neighboring products is strong and the other weak, and more profitable when the neighbors are close in quality—what we refer to as a *balanced* niche.

The preference for a balanced niche creates a novel anti-differentiation incentive in innovation. Firms will deliberately choose to innovate in a niche with a product of lower expected quality and with less horizontal differentiation if it means competition with neighboring products is balanced. This induces firms to cluster in quality not by accident but by deliberate intent.

Our second set of results turns to market dynamics and innovation over time. The ruggedness of the technological landscape generates rich dynamics, with frequent changes in direction and cycles between frontier and niche innovation, as firms search for attractive opportunities.

Driving these dynamics are two factors. The first is the randomness of outcomes obtained from innovation. How firms respond to the success or failure of a preceding innovation depends on the type of that innovation. Following frontier innovation, subsequent innovation is monotonic in outcomes. Success makes innovation more attractive, and failures cause firms to reverse course and turn to niche innovation instead. In niches, in contrast, innovation is non-monotonic in outcomes. Both successes and failures deter further innovation as they both leave behind unbalanced niches that offer little opportunity for future firms. It is only following middling outcomes that innovation may continue.

The second factor driving innovation dynamics is congestion. An innovative firm not only reveals information about product quality to future firms, it remains in the market to compete against them. As more firms exploit niches, therefore, those niches become crowded and competition becomes more intense. The red ocean becomes more red. This renders frontier innovation relatively appealing once again, and firms turn back to the blue ocean. This expands the scope of the market and, importantly, creates new niches that begin the cycle anew. In this way competition shapes innovation and innovation, in turn, shapes competition.

The market dynamics generated by the rugged landscape and congestion capture Schumpeter's (1942) process of creative destruction. In the model, some new products fail and disappear immediately, whereas others find a foothold only to be disrupted later by a new competitor, and some survive forever, retaining a place when innovation ends and the market stabilizes. Disruptive innovations can come from both frontier and niche innovation, but it is only at the frontier that a disruptive innovation induces more innovation and creative destruction is self-reinforcing. Within a niche, a breakthrough suppresses future innovation. In fact, a disruptive innovation in a niche may suppress innovation not only in the niche itself but potentially across the entire market, halting Schumpeter's process in its tracks.

Our final set of results addresses welfare. Unlike standard models, the innovativeness of the market depends not only on the number of innovations undertaken, but also on their novelty and type. Whether firms innovate boldly or incrementally, or whether they innovate in niches or expand the frontier, determines the scope and competitiveness of the market and, thus, social welfare.

An implication of this is that innovations are not necessarily all of equal social value. In our model, frontier innovations are more socially valuable than niche innovations. This is because a breakthrough at the frontier encourages further innovation whereas in a niche it suppresses innovation. As the innovative firm doesn't capture the benefit of follow-on innovation, frontier innovation is under supplied by the market. We show numerically that incentivizing frontier over niche innovation has a large positive impact on welfare.

The subtle relationship between innovation and welfare sheds new light on some old questions, in particular on how competition affects innovation, and allows us to connect the answer to welfare. In a narrow sense, competition is good for innovation in our model *a la* Arrow (1962) as it pushes firms toward more socially valuable frontier innovation and increases the novelty of their products. More broadly, however, and in line with Schumpeter (1942), competition is bad for innovation as it reduces the overall level of innovation and leads to markets that are narrower and with a sparser distribution of products. We show, nevertheless, that the impact on welfare is positive as the direct welfare benefit of more competitive markets outweighs the negative impact on innovation.[3]

Finally, we turn to the government's role in fostering innovation. On a rugged landscape, innovation flourishes at the peaks and is self-sustaining, whereas in the valleys innovation

[3]Standard treatments of this question (e.g., Aghion et al., 2005) consider only how competition affects innovation and set aside the question of welfare.

struggles and can get stuck. This implies that a policy targeted at the technological valleys can be particularly valuable. Not only can policy induce innovation where it otherwise would not occur, but as it can carry innovation across the technological valley to a peak on the other side, it has the potential to reignite self-sustaining innovation. We develop and analyze such a policy intervention and show that even when small and temporary—possibly targeting only a single firm—policy can have a large and lasting impact on innovation, with a gain in welfare larger than in smooth models of innovation (e.g., Green and Scotchmer, 1995) and that increases in the ruggedness of the technological landscape itself.

The model we develop is highly stylized. Its goal is to provide a microfoundation for innovation and competition and to link these to market outcomes. Despite the abstraction, the model connects with data and practice at several points, both in its assumptions and in equilibrium behavior. We collect and develop these touchpoints in detail in Section 6 following presentation of the formal results.

Related Literature

The Hotelling model of horizontal differentiation is a fundamental framework to understand market competition. To it we add a rugged technological landscape that is unknown to firms. 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).[4] 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

[4]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).

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.

The Hotelling framework is typically modeled with complete information. An exception is Narajabad and Watson (2011) who add a quality-ladder innovation model to a firm location model, although they restrict attention to two firms and two locations and solve the model numerically. Another line of work adds uncertainty to the Hotelling framework by assuming firms are uncertain about consumer demand (see, for example, Meagher and Zauner (2004)).

In building on the Hotelling framework, our model provides a sense of distance that captures the variable novelty and riskiness of innovation. In Letina (2016) and Bryan and Lemus (2017) firms choose the direction of their innovation, although with a finite set of directions that correspond to different projects. We allow for a continuum of correlated potential innovations. In our model, uncertainty about an innovation increases with its novelty, which 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 unbounded dimension.[5] This follows a recent search literature (Callander, 2011; Garfagnini and Strulovici, 2016; Callander and Matouschek, 2019; Carnehl and Schneider, 2024).[6] The fundamental difference with that literature is that we allow for competition. Thus, firms learn from and compete with earlier products, whereas in the search literature only information carries over from one period to the next. Competition between firms generates very different dynamics to those in Callander (2011). In Callander (2011) all frontier search precedes any niche search, and niche search iterates within the single niche where it begins. In the present model, market competition eventually drives firms out of any single niche, and innovation can cycle between frontier and niche.[7] 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.

[5]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.

[6]Jovanovic and Rob (1990) is the first use of the Brownian motion for search, limiting attention to the [0.1] interval with at most two alternatives searched.

[7]Competition also implies that optimal behavior is not Markov, as it is in the search model, in the sense that it can depend on points in the Brownian mapping beyond the nearest neighbors. This possibility arises when a neighboring product is dominated by a non-neighboring product (the neighboring product is what we refer to as inactive). In the pure search model information outside the niche is irrelevant to optimal behavior within the niche.

2 The Model

In every period $t = 1, 2, ...$ there is an infinite mass of consumers spread out evenly 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.

In the first period $t = 1$, a competitive fringe produces an initial product at location $l_0 = 0$ at zero cost, selling it at a price of zero. The product's quality is known and given by $v(l_0) > 0$. At the beginning of the period, a potential entrant, firm $t = 1$, decides whether to enter the market. If it chooses to enter, it selects a location $l_1 \in \mathcal{P}$, where it will remain thereafter. The firm incurs development costs, detailed below, but—like the competitive fringe—faces zero production costs.

The new product is an experience good, characterized by an expected quality $\mathrm{E}[v(l_1)|(l_0, v(l_0))]$, while its actual quality $v(l_1)$ becomes known only after consumption. We describe the process by which expectations are formed later. The firm then sets prices, employing third-degree price discrimination by charging different prices to consumers at different locations.

After prices are set, consumers decide whether to purchase the initial product l_0, the new product l_1, or neither. If any consumers buy the new product, its quality is revealed publicly. Finally, firms realize their profits and time advances to the next period.

In each subsequent period, $t > 1$, the competitive fringe expands to include any product introduced in period $t - 1$, which it sells at a zero price. The period then unfolds like the first: a new potential entrant, firm $t > 1$, decides whether to enter the market. If it enters, the firm selects a location $t \in \mathcal{P}$, where it will remains permanently. The expected quality of its product is given by $\mathrm{E}[v(l_t)|\mathcal{E}_t]$, where $\mathcal{E}_t$ is the set of existing products, $\mathcal{E}_t = \{(l_0, v(l_0)), (l_1, v(l_1)), ..., (l_{t-1}, v(l_{t-1}))\}$. The firm then sets prices for consumers at different locations. If any consumers purchase the product, its quality is revealed publicly. Finally, firms realize their profits and time advances to the next period.

To the Hotelling framework we add uncertainty over the quality of innovations. To capture a rugged technological landscape, we represent the mapping $v(l_j)$ from product location to product quality as the realized path of a Brownian motion with zero drift, scale $\sigma > 0$, and initial value $v(l_0)$. The firms do not know the path, and thus the quality of untried products. They do know the scale parameter σ and, thus, the expected ruggedness of the landscape.

The firms also know 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. For a frontier innovation, $l_t > l_t^f$,

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

and

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

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

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

and

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

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 red dots are the quality of the three products, as well as l_0, and the black dashed line is the expected quality of all potential products given that history. The red lines represent crudely that uncertainty over the quality of other products increases linearly on the frontier and is concave in niches.

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 consist of fixed and variable costs, where the variable cost is convex in novelty and, for simplicity, take the functional form:

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

with $c \geq \frac{5}{6\tau}$ and $F \geq 0$.[8] In a niche, new products are within the space that has already been

[8]These costs ensure that the market is always covered: if consumer $s' \in \mathcal{P}$ buys a product, all consumers

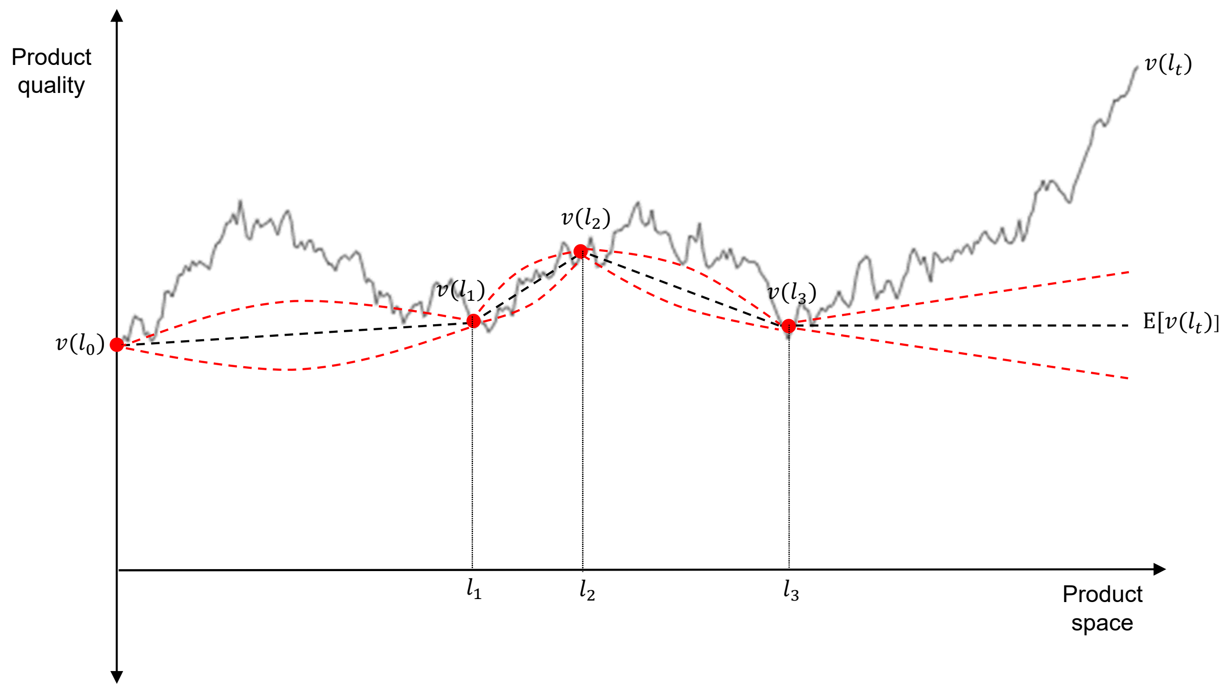

Figure 1: Beliefs on the Rugged Technological Landscape

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 continue in the market until they are no longer profitable. Additionally, we assume that the market ends in each period with a small exogenous probability, $\gamma > 0$, which plays a role only in the simulations of Section 5.

The presence of the competitive fringe ensures that entrants can only earn profits during the period they enter, as competition with the fringe drives prices and profits to zero in subsequent periods. As a result, each firm t chooses its entry, location, and pricing strategy to maximize profits in period t, taking as given the zero prices charged for all existing products, $\mathcal{E}_t$.

The equilibrium of the model is the solution to a sequence of decision-theoretic optimization problems faced by the firms. Our objective is to provide a characterization of this equilibrium both analytically and numerically.

Remark. Third-degree price discrimination and experience goods describe some products and markets but not all. The one period model of Callander and Matouschek (2022) relaxes both without substantively changing the economics of innovation and competition in this framework. Relaxing them here presents a challenge to tractability. To capture frontier inno-

$s \leq s'$ also buy a product. The assumption that markets are covered is standard in the literature on spatial competition that builds on Hotelling (1929) (see, for instance, Tirole (1988, p. 279)).

vation, a new entrant needs to be able to expand the market scope and serve new customers. This feature paired with quadratic transportation costs is intractable, though tractability is restored for linear transportation costs. Linear costs, in turn, leads to a well-known equilibrium existence problem for posted price competition, a problem that does not arise when firms can third-degree price discriminate, as we assume here. (Callander and Matouschek (2022) assume the incumbent covers the entire market; this sidesteps these problems but forecloses on the possibility of market expanding innovation.)

The assumption that the competitive fringe enters with the next innovating firm is also made for tractability, though it carries some substantive import. It reduces a firm's planning horizon to a single period, as assumed directly in other models of innovation in information rich environments (e.g., Callander, 2011). It is well known in practice that fast imitation of innovations reduces a firm's planning horizon (Klepper and Graddy, 1990). Nevertheless, imitation is not always immediate and market power can persist for some time. Exploring firms' strategy in this framework with longer horizons, and the strategic interactions this gives rise to, is of obvious interest.

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 marginal 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

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

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

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

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

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

$$\pi_t(l_t | \mathcal{E}_t) = \int_0^\infty p(s, \mathcal{E}_t) \, \mathrm{d}s - c(l_t, \mathcal{E}_t). \tag{6}$$

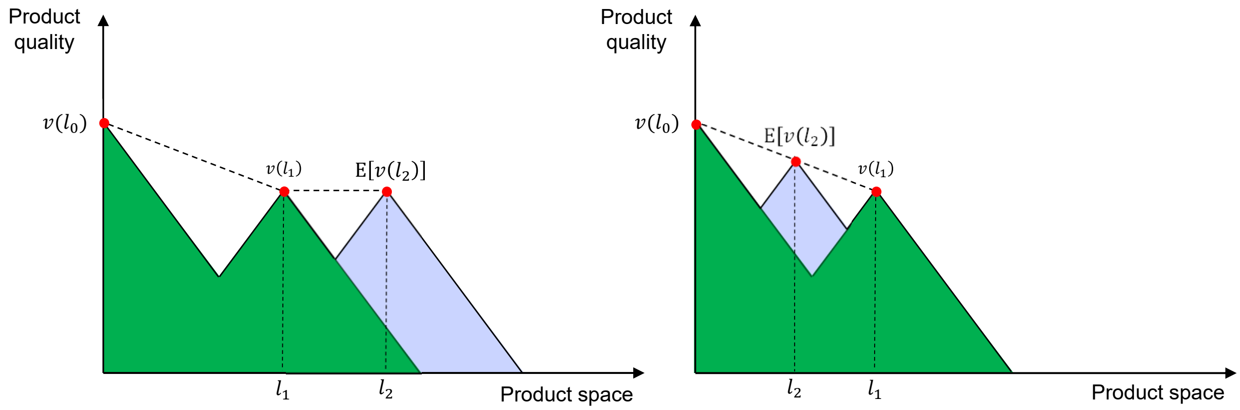

Figure 2: Value Creation & Value Capture on the Frontier (left) and in a Niche (right)

The calculation of profit represents the classic dichotomy in business strategy between *value creation* and *value capture*. The expected value a product creates is $\mathrm{E}\left[u\left(s, l_t\right)\right]$ 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, $\mathrm{E}\left[u\left(s, l_t\right)\right] - f\left(s, \mathcal{E}_t\right)$, 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 $\mathrm{E}[v\left(l_2\right)]$ 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 less R&D costs.

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 *There is a threshold* $\underline{v}_1 \geq 0$ *such that, in period one,*

(i.) if $v(l_0) \leq \underline{v}_1$*, there is no profitable innovation.*

(ii.) if $v(l_0) > \underline{v}_1$*, the optimal innovation is located at*

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

and generates profits

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

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 both created and captured by the firm. As value creation and value capture are not in conflict in this situation, the novelty of innovation is constrained only by R&D costs.[9]

Innovation at the frontier increases in the quality of the existing product, $v(l_0)$, when $v(l_0) > \underline{v}_1$ and the profit from innovation can cover the fixed cost of R&D. (Note $\underline{v}_1 = 0$ for $F = 0$.) Thus, successful products themselves induce bolder innovation. One intuition may be that a higher quality incumbent product 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, as the entering firm differentiates more, the marginal part of its competitive shadow that emerges from the incumbent's shadow is larger the higher are the product qualities.

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. This possibility represents Schumpeterian dynamics of entry-and-exit at work. Product l_1

[9]This is evident in the expression for $\pi_1(l_1^*)$ where the first term is the monopoly profit and the second term is the loss relative to that benchmark. This formulation facilitates comparison to Proposition 2 below, and readily simplifies to $\frac{\tau v(l_0)^2}{1+2c\tau} - F$.

serves customers and is active in the market when it is introduced, and it will remain active if its realized quality is not too low. However, if quality is so low that it falls within the competitive shadow of l_0, the product cannot attract customers at any price and is driven from the market. 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. The product at the frontier may have failed, but its silhouette shapes both competition and innovation going forward. 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. Using the fact that 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 threshold $\underline{v}_t \geq 0$, such that, in any period $t \geq 2$,*

(i.) if $v\left(l_t^f\right) \leq \underline{v}_t$, there is no profitable frontier innovation.

(ii.) if $v\left(l_t^f\right) > \underline{v}_t$, 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 profits

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

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

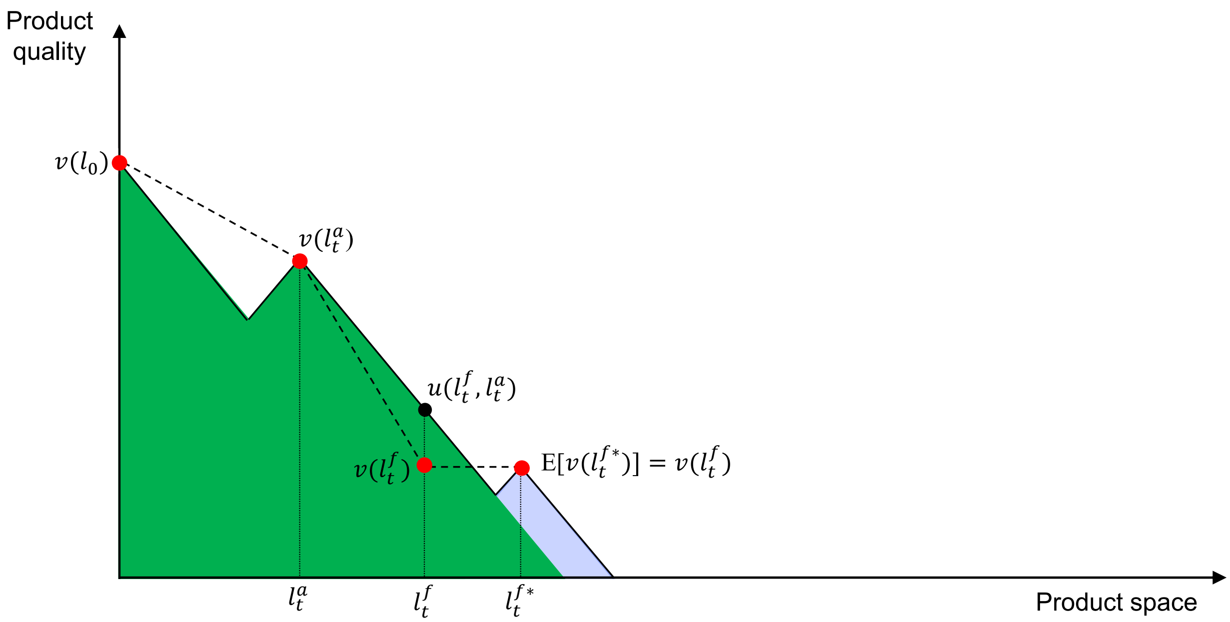

Figure 3: Frontier-Expanding Innovation with a Non-Active Frontier Product

depends on the quality of both the frontier and the largest active products, and is increasing in both. Intuitively, the threshold $\underline{v}_t$ increases and the firm's profit decreases in $u\left(l_t^f, l_t^a\right)$. For $F = 0$, there is a constant $\kappa \in (0, 1)$ such that in any period $t \geq 2$, frontier innovation is strictly profitable if and only if $v(l_t^f) > \kappa \cdot u(l_t^f, l_t^a)$. In this case, if the quality of the frontier product exceeds a fixed fraction of the shadow in which it is embedded, frontier innovation remains profitable.[10]

Proposition 2 establishes that frontier innovation can end, and that it can end even when the frontier is expected to produce positive quality products. Although it may not be in the interest of a single firm to innovate, continued frontier innovation would nevertheless raise social welfare. The reason is that firms do not internalize either the consumer surplus generated by new products or the follow-on innovations that a new frontier product may induce. These additional benefits can be sufficiently large that innovation is socially desirable even when the expected quality at the frontier is negative. This represents a market failure, one that we return to in Section 5.4.

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

[10]Setting the profit function in Proposition 2 equal to zero delivers a closed-form expression for κ, which we provide in the appendix.

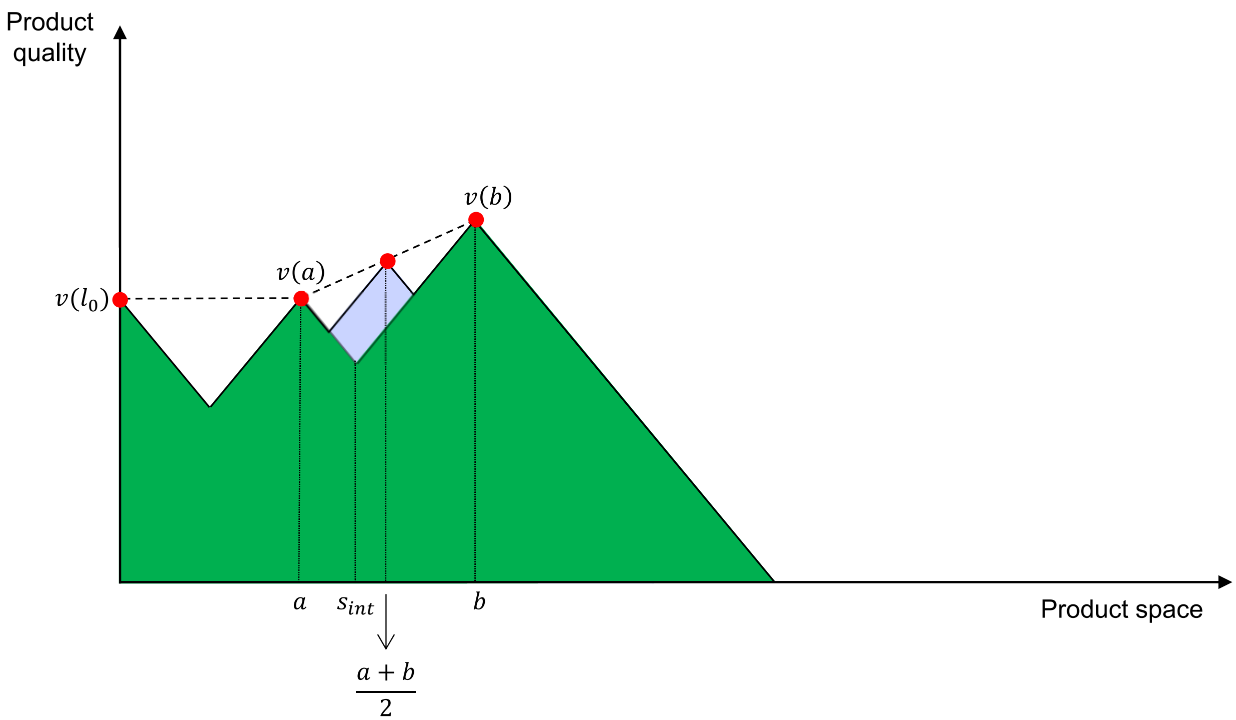

Figure 4: Niche Innovation

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 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 a *viable* niche as the set of products for which the expected quality is above the

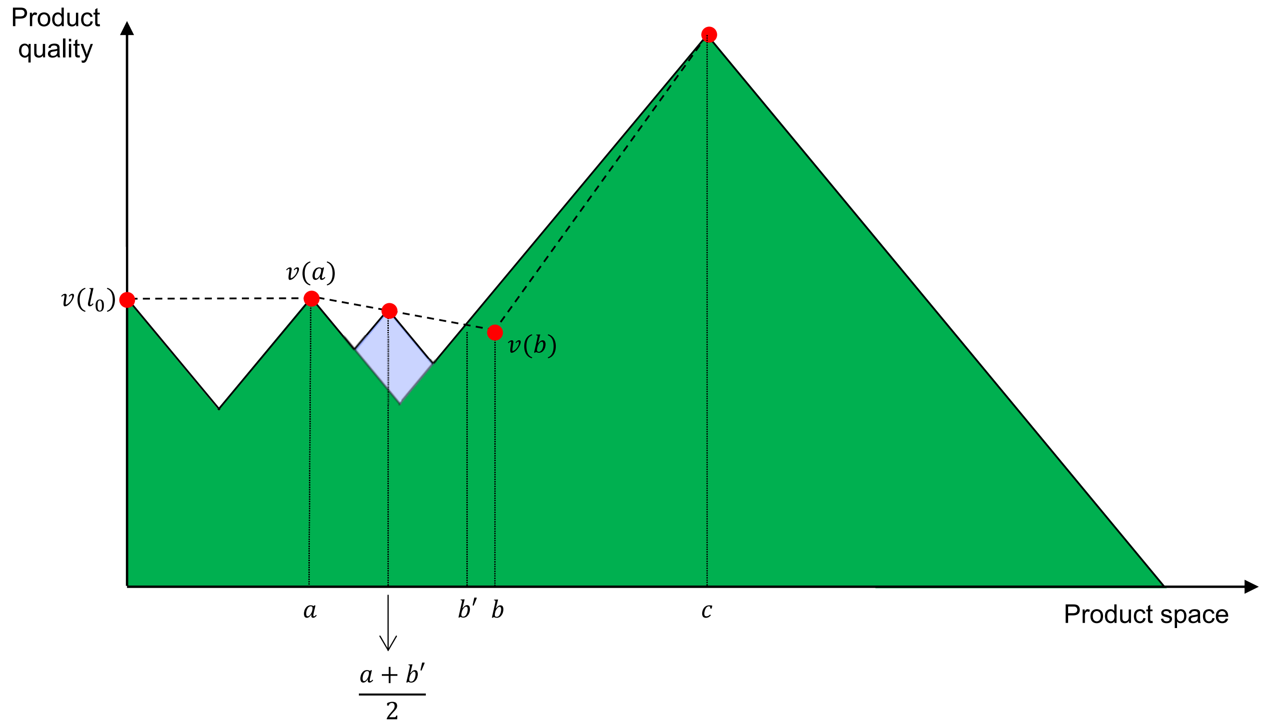

Figure 5: Niche Innovation on the Viable Sub-niche $[a, b']$

competitive shadows of other products. 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*

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

and generates profits

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

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 the midpoint of the niche 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. The entrant locates halfway along the viable niche, thereby maximizing neither competitive differentiation nor uncertainty, and competes more intensely with the neighbor that is active in the market.

The choice reflects a trade-off between market share and profitability of individual consumers. The maximum profit the entrant can make from a single consumer is by locating 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 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 the midpoint of the niche and maximally differentiates the entrant from its competitors and minimizes competition. In an *unbalanced* niche—with neighbors of unequal quality—this point is closer to the lower quality neighbor. Moving toward the higher quality neighbor lowers the profitability of a single consumer but increases the firm's market share. The entrant's market share is maximized when it locates arbitrarily close to the high quality neighbor, however this comes at the cost of intense competition. Halfway along the viable part of a niche is where these considerations 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 from Proposition 2 that $\underline{v}_t$ is the threshold below which frontier innovation is not profitable.

Proposition 4 *Suppose that frontier innovation is optimal for firm $t-1 \geq 1$. Then there exists a $\underline{v}_t$ and a $\overline{v}_t$ such that:*

(i.) if $v\left(l_t^f\right) > \overline{v}_t$, frontier innovation is optimal for firm t.

(ii.) if $v\left(l_t^f\right) \in (\underline{v}_t, \overline{v}_t]$, frontier innovation is not optimal for firm t but may be optimal for some firm $t' > t$.

(iii.) if $v\left(l_t^f\right) \leq \underline{v}_t$, 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. 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 then it disrupts the entire market, rendering all incumbent products inactive. 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.[11]

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 to exploit. 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 the viable 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.*

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. Whether the quality is high or low, all remaining products in the original niche lie in the competitive shadow of another product and are rendered inactive and driven from the market. 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

[11] In the absence of the competitive fringe, this disruptive entrant would obtain a competitive "moat" to the left due to its own success and the information gleaned from the lower quality of its predecessors.

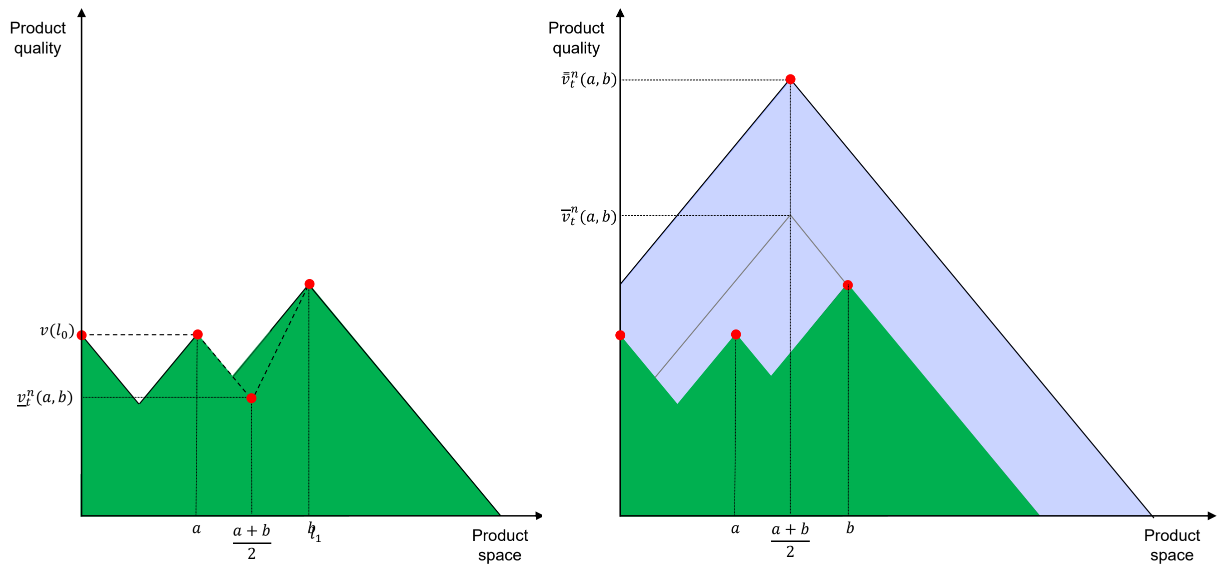

Figure 6: Entry Deterring Niche Innovation from a Failure (left) and a Success (right)

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.[12]

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.

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

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—a flat technological landscape—entry and market growth across an unbounded product space continues 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 as every niche is perfectly balanced when the technological landscape is flat. In a world of innovation, in contrast, growth is inspired by the ruggedness of the technological landscape and it is also constrained by it. The many peaks and valleys in the technological landscape 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.[13]

Three Examples. To see the richness and variety of market dynamics that are possible, it is helpful to look at three markets that exhibit contrasting patterns of innovation. These markets were generated by numerical simulations with parameters $v(l_0) = 5$, $\sigma = 1$, $\tau = 10$, $c = 0.1$, $F = 0$, and $\gamma = 0.001$.

Innovation is short-lived in the market in Figure 7. Only four products are introduced before the market stabilizes, three frontier innovations followed by a niche innovation. Of these, only one, l_2, along with product l_0, remain active when stability is reached.

The markets in Figure 8 are very different. Innovation in both of these markets is extensive and long-lasting with hundreds of new products introduced (so many that we omit marking the products in the figure; note the scale on the product space).

[13] 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.

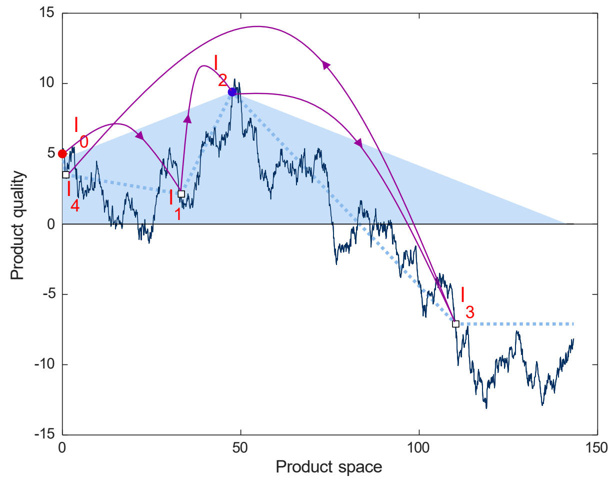

Figure 7: Short-Lived Innovation

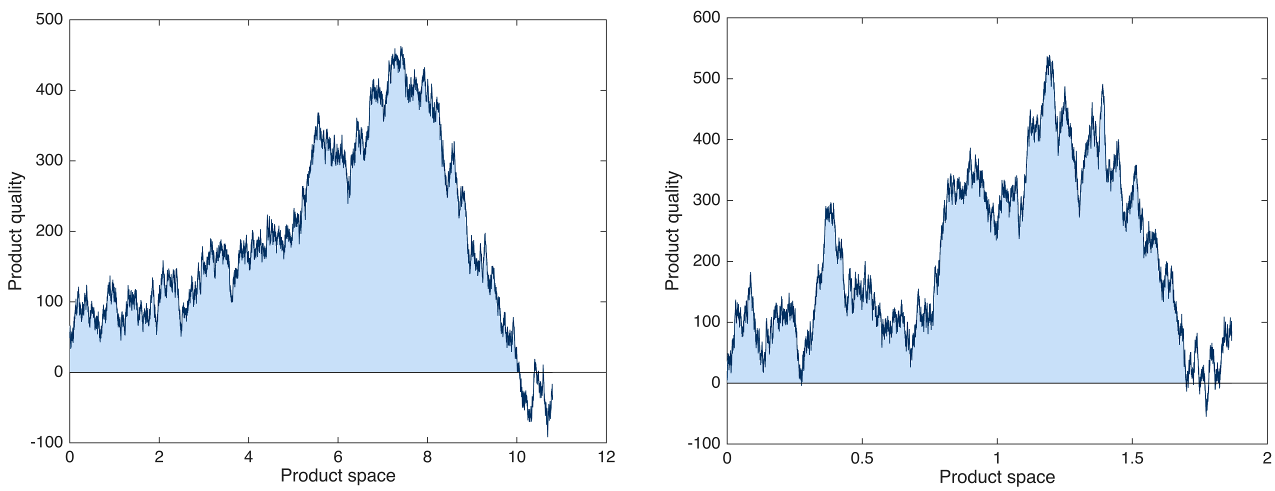

Figure 8: Long-Lived Innovation

The markets in Figure 8 also differ from each other, exhibiting contrasting growth paths. This is evident in Figure 9 that depicts growth in the size of the market—the measure of consumers served by a firm. In the left market of Figure 8, innovation is largely in two phases, beginning with a long phase of frontier innovation before switching to niche innovation as new entrants exploit the many niches that were opened up. This manifests in the left panel of Figure 9 as a market that grows in one continuous convex arc before settling down and remaining relatively stable thereafter as niche innovation plays out.

In contrast, the market on the right side of Figure 8 exhibits more changes in direction and cycles between frontier and niche innovation in three broad sweeps.[14] The cycles manifest as plateaus in the growth path, as is evident in the right side panel of Figure 9. Each increase

[14]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.

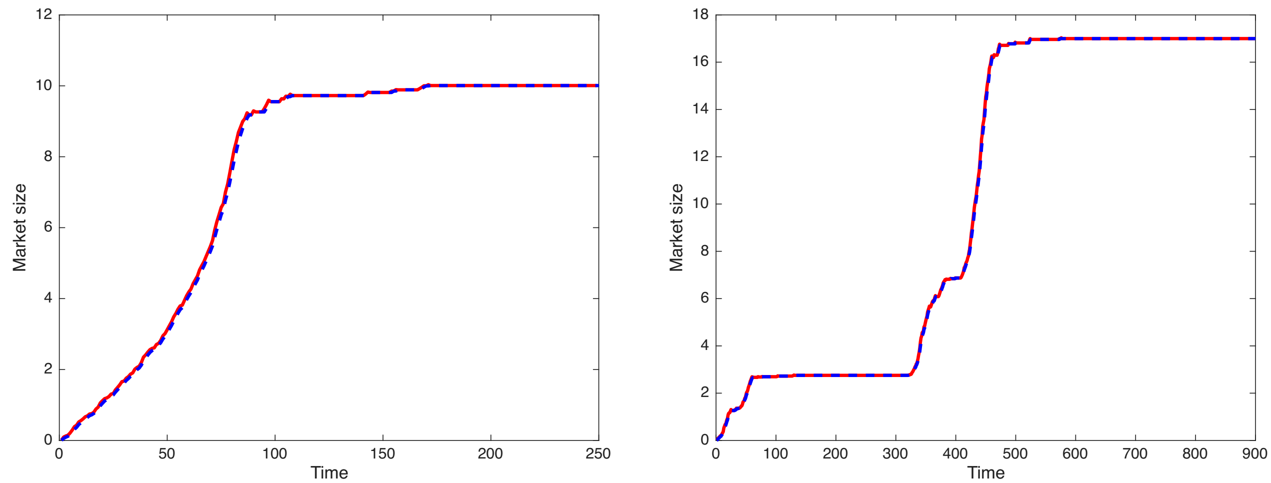

Figure 9: Market Size Growth

represents an era of frontier-expanding innovation and each plateau reflects the subsequent phase of niche innovation.[15]

5 Numerical Analysis of Market Dynamics and Welfare

5.1 Averages and Comparative Statics

To explore the range of markets and firm behavior that is possible, we ran 500,000 simulations for parameters $v(l_0) = 5$, $\sigma = 1$, $\tau = 10$, $c = 0.1$, $F = 10$, and $\gamma = 0.001$ as a benchmark case. Welfare is the sum of consumer and producer surplus, i.e., the area under all competitive shadows, less R&D costs, discounted at rate $\delta = 0.9$. Table 1 reports the results, along with the comparative statics described below.[16] Standard errors are reported in parentheses in tables and in the text.

The most prominent feature of market dynamics is a long right tail. The average number of products introduced is 41.3, although the range extends up to 7,306 products.[17] The average fraction of frontier innovation in a market is 62.9%. In aggregate, however, there is

[15]For an intuitive explanation of the difference across the two markets, observe that the left market may be aptly described as one large mountain with sub-peaks, whereas the right side market consists of several distinct mountains. A useful heuristic to predict a switch from frontier to niche innovation is the notion of *prominence* in topography—the height of a mountain's summit relative to the lowest contour line encircling it but containing no higher summit within.

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

[17]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.

more niche innovation as many markets stop early with a high fraction of frontier innovation, whereas those markets in which innovation continues for a long time see much more niche innovation. The number of cycles between frontier and niche innovation is also right skewed. The maximum number of cycles is 45, although the average number is 1.54 (0.002) with standard deviation of 1.74. In many markets, innovation is short-lived and transitions only once or a few times between innovation types, but the tail is long and some markets work through many cycles. In 95.8% of markets innovation stopped endogenously when profit opportunities were exhausted; in the remaining 4.2% of markets innovation stopped exogenously (due to the γ parameter).

To obtain further insight, we turn to comparative statics on the key model parameters.

Competition vs. Innovation: The τ parameter speaks directly to the classic question of how competition affects 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 conflicting effects on innovation.

The direct effect of an increase in competition is that, for a fixed history, innovation is less profitable. This is felt particularly acutely within market boundaries as competition cannot be avoided. Beyond the market boundaries it can, however, and the indirect effect of an increase in competition is that frontier innovation is relatively more attractive and the share of frontier innovation increases from 62.9% to 81.6%. Moreover, firms innovate more boldly at the frontier, differentiating themselves further from existing products to reduce competition, and expanding the size of the market. Consequently, frontier innovations leave behind larger niches that, indirectly, make niche innovation more attractive.

On net, the direct effect dominates the indirect effect. An increase in τ from 10 to 20 decreases the number of new products introduced from 41.3 to 14.0 and generates markets that are simultaneously narrower (serving fewer consumers) and with products that are more sparsely spaced.

Despite the dramatic decrease in innovation, the increase in competition increases welfare by 38%. This change also reflects conflicting 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 variable R&D costs also produces 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 the niches are eventually exhausted and frontier innovation does continue. When it does, it is less bold. Indirectly, therefore,

	Benchmark	$\tau = 20$	$c = 0.2$	$F = 20$	$\sigma = 2.0$
New Products	41.3	14.0	22.9	39.0	17.2
	(0.25)	(0.10)	(0.17)	(0.24)	(0.13)
Frontier Products, %	62.9%	81.6%	73.8%	61.9%	79.7%
	(0.044)	(0.034)	(0.038)	(0.045)	(0.037)
Welfare	11,495	15,917	4,087	11,099	57,747
	(77.2)	(120.1)	(25.6)	(77.7)	(547.0)

Notes. Standard errors are given in parentheses.

Table 1: Comparative statics: $v(l_0) = 5, \sigma = 1, \tau = 10, c = 0.1$, $\gamma = 0.001$, $\delta = 0.9$, $F = 10$, 500,000 simulations.

the niches left behind are narrower and less attractive, and this renders frontier innovation relatively more attractive. Nevertheless, this indirect effect is dominated by the direct effect. Despite the fraction of frontier innovations increasing from 62.9% to 73.8%, the number of frontier innovations declines along with the number of niche innovations.

This finding provides a novel twist on on how innovation is interpreted in practice. A market with high variable R&D costs has a higher fraction of frontier-expanding innovations and, therefore, may appear highly innovative. This would be misleading, however. If one were to look at the number and scale of innovation, it would become clear that innovation is fewer in number and more incremental. In our simulations, the higher propensity of frontier innovation actually represents a decrease in welfare of 64%. This result makes clear that welfare depends on the novelty and type of innovation a market produces, and not just the raw number of innovations.

A change in fixed R&D costs has a more even impact on innovation. A higher fixed cost makes frontier innovation less attractive, but does not change the novelty of the innovation that is undertaken. Thus, the size of niches that are opened are unaffected. The net impact is for slightly less innovation overall and a slightly lower fraction of frontier innovation.

Ruggedness 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 it generates more imbalanced niches. As a result, a higher proportion of innovations are at the frontier, yet even there the chance of an innovation-stopping failure is higher. An increase in σ from 1 to 2 causes the average number of new products to fall from 41.3 to 17.2. However, the chance of a breakthrough at the frontier is also higher and the innovations that are undertaken are more valuable. The breakthroughs are of sufficient quality that total welfare increases by more than 400%.[18]

5.2 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 24.1% (0.267) of products remain viable when the market stabilizes, and for those that disappear the average life-span in the market is 35.6 (0.230) periods for niche innovations and 8.8 (0.062) periods for frontier innovations.[20]

Frontier innovations not only have a shorter lifespan, they are less likely to survive through to market maturity. In the average market, frontier innovations represent 62.9% of total innovations, yet they represent only 37.6% (0.061) 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 that frontier innovations are less successful can be found in Schumpeter's creative destruction. At the frontier, success begets further innovation and creative destruction is self-reinforcing. Later firms "standing on the shoulders of giants" which is is good for society, yet 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. It stops Schumpeter's process in its tracks. 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.

[18] Arguably, R&D costs in practice are proportional to the ruggedness of the technological landscape. We consider them separately here for brevity and clarity.

[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.

5.3 Welfare and the Inefficiency of Market Innovation

It is commonly understood that the social returns to innovation exceed the private returns (Green and Scotchmer, 1995; Bloom, Schankerman and Van Reenen, 2013). This is true here as well. Our model raises the additional questions of whether the market produces the right types of innovation and whether they are sufficiently novel.

Solving for the socially optimal innovation path over a longer horizon is neither analytically possible nor computationally feasible. To gain insight into the welfare properties of market equilibrium, we turn to numerical simulation of the model and approach the question indirectly. Results reported in this section can be found in Table 2. The benchmark set of parameters is the same as in Table 1.

A taste for risk. The fundamental insight of optimal experimentation is that agents engage more risk the longer is their horizon. To explore indirectly how the myopia of firms in the model affects social welfare, we consider a variant that encourages the firms to experiment more boldly. Specifically, we endow each firm with a taste for risk that is additive to the regular profit function. The firms then optimize, and the market evolves, as before.[21]

It is illuminating to separate two cases: when a taste for risk activates at the frontier only and in a niche only. Table 2 reports the two variants. In the former case, risk at the frontier is weighted by $\alpha_f = \frac{1}{2}$ in firm utility, and in the latter case, risk in a niche is weighted by $\alpha_n = \frac{1}{2}$. In the benchmark, both parameters are zero: $\alpha_f = \alpha_n = 0$. For comparison with the benchmark parameters, welfare is reported *without* the benefit of risk included.

A taste for risk has a significant impact on search behavior and welfare when it holds for frontier innovation but not for niche innovation. At the frontier, a taste for variance encourages frontier innovation to be bolder and to continue in circumstances when it otherwise would stop. Even for a modest weight on risk, it increases the number of products searched to 44.5 from 41.3 in the benchmark, and for welfare to increase by 7.3%. The bolder innovations at the frontier opens up larger niches and this indirect effect increases the number of niche innovations in absolute terms and as a fraction of total innovations.

A taste for risk in a niche, in contrast, has no significant effect. This is not surprising given Proposition 3 for niche innovation. In a niche, the firm innovates with the product that maximizes risk if both neighboring products are active, even without an explicit taste for risk. Thus, the taste for risk in a niche only reinforces market innovation without changing it. Similarly, simulations with a general taste for risk—both in niches and at the frontier—deliver outcomes close to those with a taste for variance at the frontier alone (unreported).

These results are consistent with frontier innovation being more socially valuable than

[21] We relegate to the appendix the complete details of this variant, as well as the changes to optimal search behavior it causes. To ensure the market is always covered, a strictly positive lower bound on the fixed cost of innovation is required; $F = 10$ is sufficient.

	Benchmark	Frontier $\alpha_f = \frac{1}{2}$	Niche $\alpha_n = \frac{1}{2}$	Frontier First	Patents $p = \frac{1}{2}$
New Products	41.3	44.5	40.9	42.1	28.8
	(0.25)	(0.26)	(0.25)	(0.25)	(0.21)
Frontier Products, %	62.9%	61.4%	62.2%	63.1%	60.1%
	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)
Welfare	11,495	12,333	11,533	11,906	8,539
	(77.2)	(80.0)	(79.0)	(77.1)	(67.8)

Notes. Standard errors are given in parentheses.

Table 2: Market Interventions: $v(l_0) = 5, \sigma = 1, \tau = 10, c = 0.1$, $\gamma = 0.001$, $\delta = 0.9$, $F = 10$, 500,000 simulations.

niche innovation. Because Schumpeter's process of creative destruction is self-reinforcing at the frontier but not in niches, encouraging frontier innovation is good for welfare. The market performs better with innovation of all types, but specifically with bolder frontier innovation. In the intervention studied here this did not require a higher fraction of frontier innovation, indeed quite the opposite, although that would not seem to be an essential property. What appears to be important is that innovation at the frontier is not limited in aggregate, but that it expands the market significantly, whether that comes from a series of incremental steps or from one bold innovation.

Frontier Versus Niche Innovation. An alternative explanation of market inefficiency is that it is not the number or type of market innovations that is the problem, but rather it is the sequencing of innovation that holds back welfare. To explore this possibility, we consider an intervention that forces frontier innovation to exhaust itself before niche innovation is allowed.

This intervention, reported in Table 2 as Frontier First, has a minor but significant impact on market outcomes. The number of innovations increases from 41.3 to 42.1. The reason for this as that frontier innovation cannot be closed off by a successful niche innovation, and thus, there is more of it than in the benchmark case.

Reordering frontier and niche innovation itself seems to have no significant impact on welfare. This shows that the greater social value of frontier innovation is not in the innovations themselves but rather because creative destruction is in full force at the frontier but not in niches. A success at the frontier encourages more innovation and in that additional innovation lies the social benefit. Bringing frontier innovation forward but not changing it (beyond minimally) has, therefore, only a minor impact on welfare.

Market Power. In practice, governments often intervene in innovative markets by offering patent protection. Patents explicitly target the gap between private and social returns to welfare by granting monopoly power for a limited time to the innovative firm. In our model, the innovative firm receives the equivalent of a one-period patent—the length of time before the competitive fringe enters and competes away the value created. This is the same length of time before another innovative firm potentially enters the market with a new product, although the two lengths of time need not be equal.

To consider the effect of patents on innovation, we suppose that the length of time before the competitive fringe imitates a new product is in the control of government. Let this be a fraction of a single period of length, $p \in [0, 1]$.[22] The model until now, therefore, reflects maximum patent protection and $p < 1$ represents less protection. A firm's profit is then p times revenue less R&D costs.[23]

Shorter patent protection has, not surprisingly, a dampening effect on innovation, as evident in the final column of Table 2. Welfare decreases significantly, as does the number of innovations, and the fraction of frontier innovations. This is because R&D costs are paid regardless of patent length. This biases firms away from bold frontier innovation and instead toward more incremental frontier innovation and niche innovation where R&D costs are lower. Turning this result around, it says that patents are good for innovation and welfare, not only because they increase innovation per se, but because they push firms towards more socially valuable innovation at the frontier.[24]

5.4 Targeting the Technological Valleys

On a rugged landscape, market innovation flourishes at the peaks but struggles in the valleys. Innovation may stop in a valley even though many peaks may lie on the other side. The problem for an individual firm is that it does not know if the frontier lies in a technological valley or at the beginning of a large featureless plain, or worse yet, that the floor of the valley has not yet been reached. Even if the firm knew it was a valley, and that a peak lay on the other side, the valley may be too wide to cross in a single leap. To a society, however, the possibility of more peaks warrant the cost of crossing the valley, even if it takes many periods—and much R&D cost—to get to the other side.

Consider the situation depicted in Figure 10. In this market, innovation stops after three new products are introduced. Innovation at the frontier is deterred by the low quality for

[22] We thank an anonymous referee for this suggested formulation. It is feasible to extend to $p > 1$, although then an innovative firm must factor in competition from the next innovative firm that enters after one period of time. Delaying the arrival of that firm would avoid this difficulty, although that would be equivalent to a rescaling of time for the results reported here.

[23] Full details are provided in the appendix, along with expressions for optimal innovation.

[24] This representation captures patents crudely, though of course it leaves much out. In particular, it omits the incentive effects on innovative firms to choose the timing and novelty of subsequent innovations. We leave exploration of these possibilities to future work.

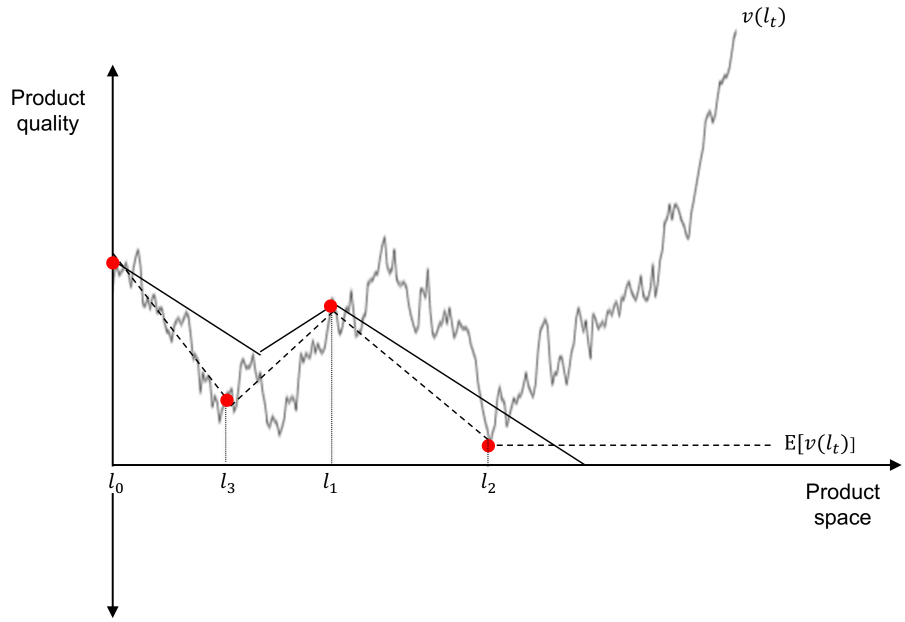

Figure 10: Incomplete Innovation

product l_2 that, whilst positive, falls too deeply within the competitive shadow of product l_1. 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 failed to improve on the quality of the initial product l_0.

Yet, as can be seen in the figure, the value created in the market represents only a small fraction of what is available. Unbeknownst to all is that the technology is highly promising, and that the failure of l_2 is only a blip on performance, albeit a particularly unlucky one. If society's knowledge could somehow cross the technological valley that l_2 lies within, innovation would restart such that even self-interested firms would be willing to innovate.

In this section, we develop and explore a policy intervention designed to help the market cross technological valleys. Specifically, we suppose that the government subsidizes the fixed cost component of R&D at the frontier when innovation has otherwise stopped. We assume the government observes the history of innovation and, thus, can identify when innovation is stuck and a subsidy is required. As only the fixed cost is subsidized, the subsidy is offered and firms accept only if the frontier product has positive quality and, thus, the product will produce positive revenue in the market.

Table 3 reports the results for three levels of fixed R&D costs and using the benchmark parameters of the previous sections otherwise. To facilitate comparison with the benchmark simulations, the subsidy in each case is equal to all but 10 of the fixed cost. As is evident in the first rows and intuition suggests, a higher fixed cost of R&D reduces the number of new products and decreases welfare.

Total Fixed Cost	30	50	90
Subsidized Fixed Cost	20	40	80
***Without* government policy:**			
New products	37.0	33.9	0
	(0.24)	(0.22)	
Welfare	10,620	9,928	125
	(75.9)	(74.3)	
***With* government policy:**			
New products (all markets)	38.9	38.6	36.8
	(0.24)	(0.24)	(0.23)
Markets with subsidy	7.14%	13.0%	100%
	(0.04)	(0.05)	
Conditional on subsidy:			
Products subsidized	1.11	1.20	1.30
	(0.002)	(0.002)	(0.0009)
Additional products induced by subsidy	25.7	28.0	35.8
	(0.74)	(0.58)	(0.23)
Welfare gain	3,508	3,903	9,654
	(130.6)	(96.5)	(73.0)
Subsidy cost	6.31	13.81	12.93
	(0.02)	(0.04)	(0.02)

Notes. Standard errors are given in parentheses.

Table 3: Policy Interventions: $v(l_0) = 5, \sigma = 1, \tau = 10, c = 0.1$, $\gamma = 0.001$, $\delta = 0.9$, 500,000 simulations.

The policy intervention creates substantial societal welfare. For a subsidy of 20, the government intervenes in only 7.14% of markets. Its interventions are short-lived, averaging 1.11 times per market and usually only for a single product, and the average cost is just 6.31 units of welfare. Yet, there are 25.7 additional products generated from the subsidy that deliver a welfare gain orders of magnitude larger than the cost of 3,508. This represents a return-on-investment of close to 55,000%. Across all markets (subsidized and unsubsidized), the subsidy represents approximately a 2.4% 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.

A subsidy of 80 provides 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 creates 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 1.30 times, with the remaining 36 odd innovations introduced because they are profitable for the firms to do so.

These results suggest that the situation in Figure 10 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.[25]

6 Empirical Connections

Although the model is highly stylized, it connects with data and practice at multiple points. We highlight here the most prominent of these touch points, both in the model's assumptions and in predicted equilibrium behavior.

The Strategic Novelty of Innovation. Two key premises of the model are (i) that innovations differ in their degree of novelty, and (ii) that firms strategically choose the novelty of their innovations. Recent empirical work that has moved beyond innovation as a yes/no decision validates both premises. In the pharmaceutical industry, Krieger, Li and Papanikolaou (2022) 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

[25] An alternative approach is for the government to subsidize research outside of the market—such as in academic or government labs—at points far from the boundaries of the existing market. If this speculative research finds a success, firms will be inspired to venture in a direction that they otherwise wouldn't. See Bryan and Williams (2021) for a thorough survey of failures in the market for innovation and possible policy interventions.

target different illnesses and segments of the patient population. Granja and Moreira (2023) provide similar evidence from financial services. Both papers also establish that the choice of novelty is strategic and deliberate. Moreover, they show that whilst firms seek out novelty for competitive reasons, they exhibit an intolerance of novelty in their innovation decisions; in Krieger, Li and Papanikolaou (2022) this behavior derives from imperfect credit markets, in our model it comes from convex costs of R&D.

Localized Technological Spillovers from R&D. Another premise of our model is that learning across products occurs by proximity—if a high quality product is found in one part of the product space, then other high quality products are likely to lie nearby. This learning represents technological spillover from R&D. That these spillovers are localized in the technological and product space in practice is a key finding of Bloom, Schankerman and Van Reenen (2013).

Localized Rivalry Spillovers from R&D. In layering innovation on top of the Hotelling line, the model ties together technological spillover with the degree of market competition. This pairing also emerges in Bloom, Schankerman and Van Reenen (2013). They show that the technological spillover from R&D benefits neighboring firms, but that these firms face a heightened competitive threat from the higher quality of the innovative firm. In Bloom, Schankerman and Van Reenen (2013) the technological spillover quantitatively dominates the rivalry spillover, and this under incentivizes investment in R&D, two facts that also emerge in our model.[26] DiMasi and Faden (2011) show for pharmaceuticals that this not only leads to heightened competition among existing nearby drugs, but that it induces, as in the model, entry into the market of new firms with nearby but not identical drugs.[27]

Profitability, Disruption, and Exit. The life cycle of firms in the equilibrium of our model matches data on product sales over time. In the model, a product space can only get more crowded, and the sales of each product (weakly) decrease monotonically over time. If a sufficiently high quality and proximate competitor emerges, sales can go to zero and the product will leave the market. This pattern of declining sales and ultimate obsolescence over a product's life-cycle is evident across a variety of products in Argente, Lee and Moreira (2024), and the entry of new and exit of old products is documented by Wollmann (2018).

Horizontal and Vertical Differentiation Over Time. Our model of horizontal and vertical innovation within a market resonates with recent empirical work on innovation within a firm. Braguinsky et al. (2021) provide evidence that as firms expand their product offerings

[26] In Bloom, Schankerman and Van Reenen's (2013) empirical specification, R&D investment improves the quality of an existing product. In our model, quality is fixed upon entry, and R&D investment is to develop and introduce a new product. Allowing for both types of R&D would be an interesting extension of our model.

[27] DiMasi and Faden (2011) also document how the speed of entry of nearby drugs into the market reduces the incentive to innovate, matching our results for patents.

horizontally, they learn about the technological and product space. In particular, they document discontinuous horizontal leaps in product characteristics followed by filling in of the newly created gaps. This matches our theoretical result of cycles between frontier and niche innovation at the market level.

The Randomness and Path Dependence of Innovation. Arguably, the key fact about innovation in practice that it is unpredictable. Klepper and Graddy (1990) provide evidence across a wide variety of new technologies how chance events and exogenous factors shape the long run evolution of an industry, both in the ultimate number and the size distribution of firms. Innovation on a rugged technological landscape captures a sense of this path dependence. In our model, industry evolution depends on the quality realization of each new product. The industry could stagnate or it could flourish depending on the realization of a single frontier innovation. Indeed, even holding constant the technological landscape, the path of an industry could be irretrievably redirected should a firm deviate in some way from its optimal product. Obviously, there are many ways that randomness can play a role in industry evolution. The framework introduced here shows how the ruggedness of the technological landscape itself can generate that randomness.

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 creation of knowledge and the practical uses of it, respectively.[28] 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).[29] 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.

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

[29]Callander (2011) offers an example of learning the drift in a pure search model.

8 Appendix: Proofs

This appendix contains proofs for results from the main model reported in Sections 3 and 4. Results for the model variants used in simulations (Section 6) are contained in the online appendix.

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

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

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

$$\pi_1(l_1) = \tau v(l_0)^2 - \frac{1}{4\tau}(2\tau v(l_0) - (l_1 - l_0))^2 - \frac{1}{2}c(l_1 - l_0)^2 - F. \tag{7}$$

Step 3: The third step is to maximize (7) with respect to location l_1. Doing so shows that the entrant generates strickly positive profits if and only if $v(l_0) > \sqrt{(1 + 2c\tau)F/\tau} \equiv \underline{v}_1$, in which case the optimal location is given by

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

and the entrant's profits at the optimal location are given by

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

If, intsead, $v\left(l_0\right)\leq\underline{v}_1$, the entrant cannot generate strickly positive profits and chooses not to innovate. ■

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,

$$l_t<l_t^a+\tau\left(v\left(l_t^a\right)-v\left(l_t^f\right)\right),\tag{8}$$

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, (8) 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

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

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 two terms are the costs of developing the entrant's product. Simplifying the above expression, we have

$$\pi_t\left(l_t\right)=\tau v\left(l_t^f\right)^2-\frac{\tau}{4}\left(v\left(l_t^a\right)+v\left(l_t^f\right)-\frac{1}{\tau}\left(l_t-l_t^a\right)\right)^2-\frac{1}{2}c\left(l_t-l_t^f\right)^2-F. \tag{9}$$

Step 4: The location that maximizes (9) is given by

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

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 $\widehat{l}_t \geq l_t^a+\tau\left(v\left(l_t^a\right)-v\left(l_t^f\right)\right)$. Instead, its revenue has to be enough to at least cover the development costs. Substituting $\widehat{l}_t$ into (9) we have

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

This expression is increasing in $v\left(l_t^f\right)$ and equal to zero for

$$v\left(l_t^f\right)>\underline{v}_t\equiv\frac{1}{3c\tau+2}\left(c\tau u\left(l_t^f,l_t^a\right)+\sqrt{2\left(2c\tau+1\right)\left(c\tau u\left(l_t^f,l_t^a\right)^2+\left(3c\tau+2\right)\frac{1}{\tau}F\right)}\right).$$

In summary, if $v\left(l_t^f\right)>\underline{v}_t$, 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, $v\left(l_t^f\right)\leq\underline{v}_t$, 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

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

and is thus given by

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

where

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

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

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

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

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

and is thus given by

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

We can then write the entrant's profits as

$$\begin{aligned}
\pi_t\left(l_t\right) &= \int_{\underline{s}_{l_t}}^{l_t} \mathrm{E}\left(v\left(l_t\right) \mid \mathcal{E}_t\right) - \frac{1}{\tau}\left(l_t - s\right)\mathrm{d}s + \int_{l_t}^{\overline{s}_{l_t}} \mathrm{E}\left(v\left(l_t\right) \mid \mathcal{E}_t\right) - \frac{1}{\tau}\left(s - l_t\right)\mathrm{d}s \\
&\quad - \int_{\underline{s}_{l_t}}^{s_{ab}} \mathrm{E}\left(v\left(a\right) \mid \mathcal{E}_t\right) - \frac{1}{\tau}\left(s - a\right)\mathrm{d}s - \int_{s_{ab}}^{\overline{s}_{l_t}} \mathrm{E}\left(v\left(b\right) \mid \mathcal{E}_t\right) - \frac{1}{\tau}\left(b - s\right)\mathrm{d}s \\
&= \frac{1}{2\tau}\left(l_t - a\right)\left(b - l_t\right)\left(1 - \left(\tau\beta\left(a,b\right)\right)^2\right).
\end{aligned} \tag{13}$$

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

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

Proof of Proposition 4: Let $\Pi_t\left(l_0, l_{t-1}^f\right)$ and $\Pi_t\left(l_{t-1}^f, l_t^f\right)$ denote the highest profit firm t can generate by locating in $\left[l_0, l_{t-1}^f\right]$ and $\left[l_{t-1}^f, l_t^f\right]$ respectively.

Suppose first that niche innovation is not profitable, that is, $\max\left[\Pi_t\left(l_0, l_{t-1}^f\right), \Pi_t\left(l_{t-1}^f, l_t^f\right)\right] = 0$. We know from Proposition 2 that frontier innovation is profitable if $v\left(l_t^f\right) > \underline{v}_t$. The threshold $\overline{v}_t$ is then given by $\overline{v}_t = \underline{v}_t$: frontier innovation is profitable, and more profitable than niche innovation, if $v\left(l_t^f\right) > \underline{v}_t$, and it is not profitable if $v\left(l_t^f\right) \leq \underline{v}_t$.

Suppose next that niche innovation is profitable, that is $\max\left[\Pi_t\left(l_0, l_{t-1}^f\right), \Pi_t\left(l_{t-1}^f, l_t^f\right)\right] > 0$. We consider two cases in turn.

Case 1 In the first case, the most profitable niche is not the right-most one,

$$\max\left[\Pi_t\left(l_0, l_{t-1}^f\right), \Pi_t\left(l_{t-1}^f, l_t^f\right)\right] = \Pi_t\left(l_0, l_{t-1}^f\right) > 0.$$

We know from Proposition 2 that frontier innovation generates profits

$$\pi_t\left(l_t^{f*}\right) = \tau v\left(l_t^f\right)^2 - \frac{c\tau}{2\left(1+2c\tau\right)}\left(v\left(l_t^f\right) + u\left(l_t^f, l_t^a\right)\right)^2 - F \geq 0 \text{ if } v\left(l_t^f\right) \geq \underline{v}_t.$$

Note that these profits are zero for $v\left(l_t^f\right) = \underline{v}_t$ and strictly increasing in $v\left(l_t^f\right)$ for all $v\left(l_t^f\right) \geq \underline{v}_t$. The threshold $\overline{v}_t$ is then given by the unique value of $v\left(l_t^f\right) > \underline{v}_t$ that solves $\pi_t\left(l_t^{f*}\right) = \Pi\left(l_0, l_{t-1}^f\right)$. Frontier innovation is (i.) profitable, and more profitable than niche innovation, if $v\left(l_t^f\right) > \overline{v}_t$, (ii.) profitable, but less profitable than niche innovation, if $v\left(l_t^f\right) \in (\underline{v}_t, \overline{v}_t]$, and (iii.) not profitable if $v\left(l_t^f\right) \in (\underline{v}_t, \overline{v}_t]$.

Finally, notice that if frontier innovation is not profitable for firm 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 may again be optimal for some firm $t' > t$. The reason is that niche innovation by firm t reduces the profits $\Pi_{t'}\left(l_0, l_{t-1}^f\right)$ that firm t' can generate by engaging in niche innovation.

Case 2 In the second case, the most profitable niche is the right-most one,

$$\max\left[\Pi_t\left(l_0, l_{t-1}^f\right), \Pi_t\left(l_{t-1}^f, l_t^f\right)\right] = \Pi_t\left(l_{t-1}^f, l_t^f\right) > 0.$$

Since $\Pi_t\left(l_{t-1}^f, l_t^f\right) > 0$, it must be that

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

If the first inequality did not hold, consumers in $\left[l_{t-1}^f, l_t^f\right]$ would prefer l_{t-1}^a to any product in $\left[l_{t-1}^f, l_t^f\right]$, and if the second inequality did not hold, they would prefer l_t^f to any product in

$\left[l_{t-1}^f, l_t^f\right]$. For what follows, it is useful to rewrite (14) as

$$L \leq v\left(l_t^f\right) \leq H, \tag{15}$$

where $L \equiv u\left(l_{t-1}^f, l_{t-1}^a\right) - \frac{1}{\tau}\left(l_t^f - l_{t-1}^f\right)$ and $H \equiv v\left(l_{t-1}^f\right) + \frac{1}{\tau}\left(l_t^f - l_{t-1}^f\right)$. The right-most niche is then given by

$$\left[l_{t-1}^f + \tau \frac{u\left(l_{t-1}^f, l_{t-1}^a\right) - v\left(l_{t-1}^f\right)}{1 + \tau\beta\left(l_{t-1}^f, l_t^f\right)}, l_t^f\right],$$

where

$$\beta\left(l_{t-1}^f, l_t^f\right) = \frac{v\left(l_t^f\right) - v\left(l_{t-1}^f\right)}{l_t^f - l_{t-1}^f} \tag{16}$$

and

$$l_t^f = l_{t-1}^{f*} = l_{t-1}^f + \frac{\tau}{1 + 2c\tau}\left(v\left(l_{t-1}^f\right) + u\left(l_{t-1}^f, l_{t-1}^a\right)\right). \tag{17}$$

Next, we use the expressions in Propositions 2 and 3, to express the profits firm t generates by engaging in frontier innovation minus the profits it generates by innovating in the niche as

$$\begin{aligned}\Delta\left(v\left(l_t^f\right)\right) \equiv\ & \tau\frac{v\left(l_t^f\right)^2}{1 + 2c\tau} - F \\ & -\frac{1}{8\tau}\left(l_t^f - l_{t-1}^f - \tau\frac{u\left(l_{t-1}^f, l_{t-1}^a\right) - v\left(l_{t-1}^f\right)}{1 + \tau\beta\left(l_{t-1}^f, l_t^f\right)}\right)^2\left(1 - \tau^2\beta\left(l_{t-1}^f, l_t^f\right)^2\right),\end{aligned}$$

where $\beta\left(l_{t-1}^f, l_t^f\right)$ and l_t^f are again given by (16) and (17). It is useful to rewrite this expression as

$$\Delta\left(v\left(l_t^f\right)\right) = \left(\frac{v\left(l_t^f\right) - v\left(l_{t-1}^f\right)}{l_t^f - l_{t-1}^f} - \left(-\frac{1}{\tau}\right)\right)^{-1} f\left(v\left(l_t^f\right)\right), \tag{18}$$

where

$$f\left(v\left(l_t^f\right)\right) = \left(\tau\frac{v\left(l_t^f\right)^2}{1+2c\tau} - F\right)\left(\frac{v\left(l_t^f\right) - v\left(l_{t-1}^f\right)}{l_t^f - l_{t-1}^f} - \left(-\frac{1}{\tau}\right)\right)$$
$$-\frac{\tau}{8}\left(l_t^f - l_{t-1}^f\right)^2\left(\frac{1}{\tau} - \frac{v\left(l_t^f\right) - v\left(l_{t-1}^f\right)}{l_t^f - l_{t-1}^f}\right)\times$$
$$\times\left(\frac{v\left(l_t^f\right) - v\left(l_{t-1}^f\right)}{l_t^f - l_{t-1}^f} - \left(-\frac{1}{\tau} + \frac{u\left(l_{t-1}^f, l_{t-1}^a\right) - v\left(l_{t-1}^f\right)}{l_t^f - l_{t-1}^f}\right)\right)^2.$$

We are interested in the sign of $\Delta\left(v\left(l_t^f\right)\right)$ for any $v\left(l_t^f\right) \in [L, H]$. Since the first term on the right-hand side of (18) is positive, the sign of $\Delta\left(v\left(l_t^f\right)\right)$ is the same as the sign of $f\left(v\left(l_t^f\right)\right)$. We proceed in three steps.

The first step is to note that

$$\text{sign}\Delta\left(v\left(l_t^f\right)\right) = \text{sign}f\left(v\left(l_t^f\right)\right) = \text{sign}\left(\tau\frac{v\left(l_t^f\right)^2}{1+2c\tau} - F\right) \text{ if } v\left(l_t^f\right) \in \{L, H\},$$

where the last term are the profits firm t makes if it engages in frontier innovation. If $v\left(l_t^f\right) \in \{L, H\}$, profits from niche innovation are zero, so that the sign of $\Delta\left(v\left(l_t^f\right)\right)$ is positive if frontier innovation is profitable and negative otherwise.

The second step is to evaluate the $f'(L)$. Differentiating $f\left(v\left(l_t^f\right)\right)$ we have

$$\begin{aligned} f'\left(v\left(l_t^f\right)\right) &= \frac{2\tau v_b}{1+2c\tau}\left(\frac{v\left(l_t^f\right)-v\left(l_{t-1}^f\right)}{l_t^f-l_{t-1}^f}-\left(-\frac{1}{\tau}\right)\right) \\ &+\frac{\tau\left(l_t^f-l_{t-1}^f\right)}{8}\left(\frac{v\left(l_t^f\right)-v\left(l_{t-1}^f\right)}{l_t^f-l_{t-1}^f}-\left(-\frac{1}{\tau}+\frac{u\left(l_{t-1}^f,l_{t-1}^a\right)-v\left(l_{t-1}^f\right)}{l_t^f-l_{t-1}^f}\right)\right)^2 \\ &-\frac{\tau\left(l_t^f-l_{t-1}^f\right)}{4}\left(\frac{1}{\tau}-\frac{v\left(l_t^f\right)-v\left(l_{t-1}^f\right)}{l_t^f-l_{t-1}^f}\right)\times \\ &\times\left(\frac{v\left(l_t^f\right)-v\left(l_{t-1}^f\right)}{l_t^f-l_{t-1}^f}-\left(-\frac{1}{\tau}+\frac{u\left(l_{t-1}^f,l_{t-1}^a\right)-v\left(l_{t-1}^f\right)}{l_t^f-l_{t-1}^f}\right)\right) \\ &+\frac{1}{l_t^f-l_{t-1}^f}\left(\frac{\tau v\left(l_t^f\right)^2}{1+2c\tau}-F\right). \end{aligned}$$

Substituting $v\left(l_t^f\right)=L$ we obtain

$$f'(L)=\frac{2\tau L\left(u\left(l_{t-1}^f,l_{t-1}^a\right)-v\left(l_{t-1}^f\right)\right)}{(1+2c\tau)\left(l_t^f-l_{t-1}^f\right)}+\frac{1}{l_t^f-l_{t-1}^f}\left(\frac{\tau L^2}{1+2c\tau}-F\right),$$

where the first term on the right-hand side is positive and the sign of the second is the same as the sign of the term in brackets, which is the profit firm t would make if it engages in frontier innovation. We, therefore, have that $f'(L)>0$ if frontier innovation is profitable when $v\left(l_t^f\right)=L$.

The third step is to show that $f\left(v\left(l_t^f\right)\right)$ is convex for all $v\left(l_t^f\right)\in[L,H]$. Notice first that $f\left(v\left(l_t^f\right)\right)$ is a third-order polynomial that is concave if $v\left(l_t^f\right)$ is below a threshold and convex if it is above. To show that $f\left(v\left(l_t^f\right)\right)$ is convex for all $v\left(l_t^f\right)\in[L,H]$ it is, therefore, sufficient to show that $f''(L) > 0$. To this end, note that

$$\begin{aligned} f''\left(v\left(l_t^f\right)\right) &= \frac{2\tau}{1+2c\tau}\left(\frac{v\left(l_t^f\right)-v\left(l_{t-1}^f\right)}{l_t^f-l_{t-1}^f}-\left(-\frac{1}{\tau}\right)\right) \\ &+\frac{4\tau v\left(l_t^f\right)}{(1+2c\tau)\left(l_t^f-l_{t-1}^f\right)} \\ &+\frac{\tau}{2}\left(\frac{v\left(l_t^f\right)-v\left(l_{t-1}^f\right)}{l_t^f-l_{t-1}^f}-\left(-\frac{1}{\tau}+\frac{u\left(l_{t-1}^f,l_{t-1}^a\right)-v\left(l_{t-1}^f\right)}{l_t^f-l_{t-1}^f}\right)\right) \\ &-\frac{\tau}{4}\left(\frac{1}{\tau}-\frac{v\left(l_t^f\right)-v\left(l_{t-1}^f\right)}{l_t^f-l_{t-1}^f}\right). \end{aligned}$$

Substituting $v\left(l_t^f\right)=L$ we obtain

$$f''(L)=\frac{1}{4}\frac{4(6c\tau-5)\,v\left(l_{t-1}^f\right)+(4c\tau(c\tau+12)+7)\left(u\left(l_{t-1}^f,l_{t-1}^a\right)-v\left(l_{t-1}^f\right)\right)}{\left(u\left(l_{t-1}^f,l_{t-1}^a\right)+v\left(l_{t-1}^f\right)\right)(2c\tau+1)}>0,$$

where the sign follows from $c\tau \geq 1$.

We can now sign $f\left(v\left(l_t^f\right)\right)$, and thus $\Delta\left(v\left(l_t^f\right)\right)$. Suppose first that frontier innovation is profitable, that is,

$$\frac{\tau v\left(l_t^f\right)^2}{1+2c\tau}-F>0.$$

From the above we know that in this case

$$f(L)>0,\ f'(L)>0 \text{ and } f''\left(v\left(l_t^f\right)\right)>0 \text{ for all } v\left(l_t^f\right)\in[L,H].$$

The threshold $\overline{v}_t$ is then given by $\overline{v}_t=\underline{v}_t$: frontier innovation is profitable, and more profitable than niche innovation, if $v\left(l_t^f\right)>\underline{v}_t$, and it is not profitable if $v\left(l_t^f\right)\leq\underline{v}_t$.

Suppose next that frontier innovation is not profitable. From the above we know that in this case

$$f(L)\leq 0,\ f(H)>0 \text{ and and } f''\left(v\left(l_t^f\right)\right)>0 \text{ for all } v\left(l_t^f\right)\in[L,H].$$

The threshold $\overline{v}_t$ is then given by the unique value of $v\left(l_t^f\right)\geq\underline{v}_t$ for which $f\left(v\left(l_t^f\right)\right)$, and thus $\Delta\left(v\left(l_t^f\right)\right)$, are zero. Frontier innovation is (i.) profitable, and more profitable than niche innovation, if $v\left(l_t^f\right) > \overline{v}_t$, (ii.) profitable, but less profitable than niche innovation, if $v\left(l_t^f\right) \in (\underline{v}_t, \overline{v}_t]$, and (iii.) not profitable if $v\left(l_t^f\right) \in (\underline{v}_t, \overline{v}_t]$.

Finally, as we noted above, if frontier innovation is not profitable for firm 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 may again be optimal for some firm $t' > t$. The reason is that niche innovation by firm t reduces the profits $\Pi_{t'}\left(l_0, l_{t-1}^f\right)$ 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 is 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 is no longer viable from one period to the next is bounded below by:

$$CDF\left(-\frac{x}{\tau}\right) + \left(1 - CDF\left(\frac{x}{\tau}\right)\right) = 1 - \text{erf}\left(\frac{\sqrt{2x}}{\tau}\right), \tag{19}$$

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 shorter than in the hypothetical, and (19) 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

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.

Argente, David, Munseob Lee and Sara Moreira. 2024. "The life cycle of products: Evidence and implications." *Journal of Political Economy* 132(2):337–390.

Arrow, Kenneth. 1962. Economic Welfare and the Allocation of Resources for Invention. In *The Rate and Direction of Inventive Activity: Economic and Social Factors*. Princeton, NJ: Princeton University Press pp. 609–626.

Bloom, Nicholas, Mark Schankerman and John Van Reenen. 2013. "Identifying technology spillovers and product market rivalry." *Econometrica* 81(4):1347–1393.

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 Heidi L. Williams. 2021. Innovation: market failures and public policies. In *Handbook of industrial organization*. Vol. 5 Elsevier pp. 281–388. Issue 1.

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.

Carnehl, Christoph and Johannes Schneider. 2024. "A Quest for Knowledge.".

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.

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. 2023. "Product innovation and credit market disruptions." *The Review of Financial Studies* 36(5):1930–1969.

Green, Jerry R and Suzanne Scotchmer. 1995. "On the division of profit in sequential innovation." *The Rand journal of economics* pp. 20–33.

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.

Jovanovic, Boyan and Rafael Rob. 1990. "Long Waves and Short Waves: Growth Through Intensive and Extensive Search." *Econometrica* 58(6):1391–1409.

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

Klepper, Steven and Elizabeth Graddy. 1990. "The evolution of new industries and the determinants of market structure." *The RAND Journal of Economics* pp. 27–44.

Krieger, Joshua, Danielle Li and Dimitris Papanikolaou. 2022. "Missing Novelty in Drug Development." *The Review of Financial Studies* 35(2):636–679.

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.

Narajabad, Borghan and Randal Watson. 2011. "The dynamics of innovation and horizontal differentiation." *Journal of Economic Dynamics and Control* 35(6):825–842.

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.

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.

Tirole, Jean. 1988. *The theory of industrial organization*. MIT press.

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