

SEARCH, BARGAINING, AND EXPERIMENTATION

Arrow Meets Hotelling: Modeling Spatial Innovation[†]

By STEVEN CALLANDER, NICOLAS LAMBERT, AND NIKO MATOUSCHEK*

Any theory that purports to explain novelty, whether it deal with invention, innovation, or the emergence of new species of biota, is intrinsically difficult and paradoxical. How can you have a theory of the unexpected? If you can understand what novelties will emerge, they would not be novelties. (Arrow 2012, p. 43)

Despite a late and slow start in economics, innovation has now moved to its rightful place at the center of the field. Innovation drives economic growth and development; without it, progress would be incremental at best. An important objective of the study of economics, therefore, is to obtain a model, or set of models, that capture the fundamental ingredients of innovation. As the opening quote from Kenneth Arrow lays out, this is a thorny problem.

Innovation is difficult to pin down because it is not just a matter of doing things better, but it is also a matter of doing things differently. Sometimes these differences are small, and sometimes they are enormous, with ideas and technologies that depart radically from those that came before. The iPhone was such a breakthrough over the Blackberry because it not only did mobile communication better, but it did it differently. As Arrow (2012, p. 44) continued, “it is precisely the way the new ... innovation differs from the present that is of interest, and that is what is difficult to predict.”

* Callander: Stanford GSB, Knight Management Center (email: sjc@stanford.edu); Lambert: Department of Economics, MIT (email: nicolas.lambert@mit.edu); Matouschek: Kellogg School of Management, Northwestern University (email: n-matouschek@kellogg.northwestern.edu).

[†] Go to <https://doi.org/10.1257/pandp.20211051> to visit the article page for additional materials and author disclosure statement(s).

The existing approaches in the literature suppress these differences. Innovation is formulated as a choice of whether but not how. Firms decide whether and how much to invest in innovation and, with luck, are successful. This allows them to produce the *same* product at lower cost or with higher quality, as in Arrow’s (1962) original model and “quality ladder” models (Reinganum 1982, Aghion et al. 2005). Or the innovation produces a new product variety, as in Romer’s (1990) model that employs a Dixit-Stiglitz–style model of monopolistic competition in which products are symmetrically different.

In practice, innovation is more than mere investment. Firms choose *how* they want to innovate and not only *whether* to innovate.¹ An innovator can choose to be incremental, barely tweaking existing products, or she can be bold, developing a product that differs radically from those that came before. Apple could have innovated incrementally on the Blackberry, improving the interface or simply combining a phone with their iPod music player. Instead Apple chose to innovate radically.

The choice of novelty—the extent to which a product is horizontally differentiated from those that came before—is important for both innovation and competition. It is important for innovation because of the risk it brings. More radical innovations come with higher risk. In choosing the iPhone, Apple chose a high risk path. Would the new product—without a physical keyboard—work technologically? Would consumers like the different design? A more incremental product would not have faced the same risk, although it also would not have offered the same potential rewards. This is why firms

¹ Acemoglu’s (1998) model of directed technical change takes a step in this direction, adding to Romer (1990) the choice of whether an innovation is labor or capital augmenting.

that innovate boldly are more likely to produce breakthrough successes but also more likely, when they fail, to do so spectacularly.

The novelty of an innovation is important for competition because in choosing how radically to innovate, a firm is also choosing how much to compete against existing products. The bolder it innovates, the more different is the set of consumers it targets. Incremental innovations compete for the existing market; bold innovations can open up new segments of the market.

In this note we describe an approach that formalizes these ideas. We model the novelty of an innovation—from incremental to bold and everything in between—in a way that captures differences in risk and show how to combine this with a model of competition in markets. The approach represents a meeting of Arrow and Hotelling, marrying the uncertainty and novelty of innovation with spatial competition and product differentiation. Formally, we meld two off-the-shelf tools, combining Hotelling's (1929) competition model with the model of search and experimentation in Callander (2011). The unifying element is the importance of distance, to competition in the former and to the novelty and uncertainty of innovation in the latter. We report on and extend two companion papers that apply the model to classic questions of innovation and competition.

I. Novelty and Risk

Innovation may be uncertain, but it need not be blind. Innovation in practice is *directed*, and an innovator relies on two sources of available information. She has theoretical knowledge about the underlying technology as well as the experience of those who came before. Her theoretical knowledge is premised on the fact that the quality of a product is related to products nearby in the technology space. Nearby products are likely to produce nearby outcomes. The innovator can use this knowledge to predict the quality of untried products.

That the quality of products is related is also why experience can guide innovation. Knowledge of one product's value provides an anchor from which expectations about other products can be formed. An innovator combines her theoretical knowledge with practical experience to make predictions about which products are likely to succeed. To be sure, the accuracy

of her predictions diminish in how far a product is from what is known concretely, which is only to say that the risk of an innovation increases in how bold and how novel it is.

As intuitive as this description of the innovator's problem is, it is missing from existing models. In the dominant models of innovation, as well as standards models of experimentation and search, innovation is not directed. Innovators either do not choose a new product at all or they choose one randomly.

To make the sense of space and distance between products meaningful, and to capture a sense of *directed* innovation, a literature has emerged using stochastic processes to represent the mapping from products to quality. Formally, firms face a continuum of potential products to develop, represented by the real line, \mathbb{R} . Only one of these products has previously been tried, the product at 0, which was shown to have quality v_0 . The quality of all other potential products is unknown until they are tried.

A firm's theoretical knowledge is how the mapping from products to quality is generated. This is where stochastic processes become useful in capturing the uncertainty of innovation while providing a theoretical framework. In this regard, a particularly tractable process is the Brownian motion with drift μ and variance σ^2 . This means that for the product at location l_1 , the quality is a draw from a normal distribution with mean $v_0 + \mu l_1$ and variance $\sigma^2 |l_1|$. This captures the intuitive property that uncertainty about the quality of a new product increases in the distance from a known product. The drift parameter represents the underlying technology, representing the expected change in quality as we move through the product space.

Innovation is often time consuming and can incur substantial development costs. These costs increase in the novelty of an innovation, reflecting the fact that it takes time and resources to research into the unknown. When applicable, let these costs be given by $c(l) \geq 0$, with $c(0) = c'(0) = 0$, $c'(l) \geq 0$, and $c''(l) \geq 0$.

Meeting Hotelling.—We marry this model of innovation with Hotelling's (1929) model of market competition. Consumers have preferences over different products, with ideal points distributed uniformly over the product space. The utility from a product at location l decreases in the distance of that product from

the consumer's ideal. Consumers also care about the realized quality of a product as well as, of course, the price. The reservation utility of buying neither product is zero.

Market competition can take many forms. The following simple structure allows us to address specific questions about innovation and competition, although the possibilities are many. An incumbent firm is present in the market, producing the only previously tried product located at $l_0 = 0$ and known to have quality v_0 . In each period $t = 1, 2, \dots$, a single innovator enters the market by choosing a location $l_t \in \mathbb{R}$. The firms in the market then simultaneously set prices.

We endow firms with the ability to third-degree price discriminate; thus, a pricing strategy is a function giving a price for each consumer type conditional on the quality of the two products.

For simplicity, let each consumer, s , buy either zero or one unit of a product; set the marginal cost of production to zero for all firms; and assume that the products are experience goods—such that consumers buy before the realization of quality in the period a product is introduced. We characterize the unique subgame perfect equilibrium for all permutations of the model. To conserve space, we omit notation unless necessary.

II. The Arrow Replacement Effect

In his famous discourse on innovation, Arrow (1962) posited a specific claim about competition and innovation. This claim, now known as the Arrow replacement effect, argues that the incentive to innovate is higher for a duopolist than it is for a monopolist. A monopolist, Arrow observed, only gains the marginal benefit of an innovation, whereas a duopolist also steals market share from its competitor. This additional benefit of innovation is the extra impetus for a competitive firm to innovate.

Callander and Matouschek (2020) examine the relative incentive to innovate when innovators choose the novelty rather than the intensity of their innovation. Using a one-period version of the model, they ask, is innovation bolder if the entrant is independent and must compete with the incumbent, or if the entrant is owned by the incumbent itself? This offers a new perspective on Arrow's question. This perspective matters because the novelty of innovation is tied so closely to risk. If competition induces bolder

innovation, then it is from competition that we can expect breakthrough innovations to come (as well as the more spectacular failures).

Innovation and competition interact in different ways in terms of novelty. In this setting, the market-stealing logic of Arrow is flipped. The entrant steals more market share the closer it locates to the incumbent and the less innovative is its product. The incentive to steal market share, therefore, pushes an independent entrant toward a more incremental innovation than an incumbent-owned entrant, who would only cannibalize its own sales.

Countering this is the classic Hotelling incentive for differentiation. The closer two products are, the more intense is price competition between them. The independent entrant has an incentive to innovate more boldly because doing so softens competition with the incumbent. The monopolist, who can coordinate prices across its products, does not have this incentive.

Callander and Matouschek (2020) show that without price discrimination, the incentive to soften price competition outweighs the desire to steal market share, and in equilibrium a duopolist innovates more boldly than does the incumbent. This aligns with the ordering in Arrow's original claim, although the reason is now despite, not because of, the incentive to steal market share. They refer to this as the *spatial Arrow replacement effect*.

To see these effects more clearly, consider a simplified version of their model in which firms are able to third-degree price discriminate. In this extreme case, the market-stealing and competition-softening incentives exactly balance out. Consequently, the entrant innovates to exactly the same degree whether it is independent or controlled by the incumbent.

PROPOSITION 1: *The independent entrant and the incumbent innovate to the same degree.*

Figure 1 depicts the case in which the quality of the innovation is also v_0 for an incumbent-owned entrant (left panel) and an independent entrant (center panel). The market-stealing and competition effects can be seen in the competitive shadow that each product casts. Consumers common to both shadows are willing to buy either product. The monopolist has an incentive to minimize this region so as to reduce cannibalization and increase the market




FIGURE 1. SPATIAL ARROW EFFECT (i) INCUMBENT OWNED, (ii) INDEPENDENT ENTRANT, (iii) WITH SWITCHING COSTS

span of its firm. The independent entrant also wants to reduce the size of the overlap because competition is more intense when the consumer has a choice.

Proposition 1 implies that these incentives are exactly equal. To see why, note that perfect price discrimination implies that a firm is able to capture all of the *marginal* value it creates. This does not mean profit is independent of the entrant’s ownership, but it does imply that the incentive to innovate is the same.

This can be seen in market pricing in each situation. A monopolist sells to each consumer the product that delivers the highest value, charging a price equal to that value. To sustain this high price when the shadows overlap, the incumbent does not offer the lower-valued product for sale (or prices unreasonably high). The incumbent’s profit is the total area under both shadows, and consumer surplus is zero. Critically, the marginal gain from innovation is the additional area created under the new product’s shadow.

The independent entrant doesn’t have the luxury of coordinating prices with the incumbent. When the shadows overlap, competition drives the price of the less valued product to zero. The consumer still buys the higher-valued product but now only pays the marginal difference in qualities. Outside of the overlap, each firm has total market power and charges a price equal to the consumer’s value. Industry profit is lower in this case, and consumer surplus is positive, as indicated in the figure. Critically, however, the entrant’s profit—the marginal gain from innovation—is again the additional area created under the new product’s shadow. This generates the equivalence of Proposition 1.

This knife-edged balance is only between the forces of price competition and market stealing, and the presence of other forces will restore the spatial Arrow replacement effect.

For instance, an independent entrant typically faces the additional need to convince consumers to switch from the incumbent product. An incumbent-owned entrant can smooth the transition through the use of a common brand. This can be formalized by imposing a cost on a consumer who switches to the independent entrant’s product, with zero (or a lesser) cost if switching is within products of the same brand. Adding switching costs restores the spatial Arrow effect.

PROPOSITION 2: *Switching costs restore the spatial Arrow replacement effect.*

This is depicted in the right-side panel of Figure 1. The entrant wins a smaller market share and can only charge a lower price to those who switch. This lowers the entrant’s profit. More importantly, moving further to the right reduces the size of this loss and, thus, the independent entrant has an additional incentive to innovate boldly to appeal more to new consumers.

III. Directed Innovation

The most important questions about innovation are arguably about dynamics. Does a breakthrough beget further breakthroughs? Or is the feedback loop more muted, mean reverting, or even negative? The framework we introduce offers a different perspective on these questions. It allows us to ask not only whether innovation persists over time but also how the *type*, *direction*, and *novelty* of innovation evolve.

We explore these questions in Callander, Lambert, and Matouschek (2020) using the dynamic version of the model with zero drift and a sequence of independent entrants. We adopt the convention that market power from an innovation is short lived and suppose that the advantage is only for a single period. For

simplicity, assume development costs are negligible.

The first entrant, as we saw above, chooses only the novelty of her innovation. Subsequent entrants face an additional choice. With multiple products already in the market, an entrant has multiple regions in which she can innovate. One option is to locate between existing products, to try to win over consumers to a targeted product. We refer to this as *niche* innovation. Alternatively, the entrant can explore new ground and locate beyond the boundaries of existing products, expanding the set of consumers served in the market. We refer to this as *frontier-expanding* innovation. The entrant chooses the type and direction of her innovation as well as its novelty. The distinction between niche and frontier-expanding innovation provides a foundation for Kim and Mauborgne's (2004) popular management notion of red ocean and blue ocean strategies, respectively.

Niche innovation is an exhaustible resource. The more firms pursue niches, the less lucrative they are to exploit, and the more crowded they become. Frontier-expanding innovations open up new niches, expanding the possibility for continued innovation. This generates cycles between niche and frontier-expanding innovation. Entrants repeatedly return to the frontier not because of any new insight but rather simply because the existing field of competition has become too crowded. The cycles between the types of innovation are highly irregular, as the vicissitudes of innovation make the relative appeal of each strategy subject to randomness.

Innovation of both types eventually ends, and the market stabilizes. It ends at the frontier when a new product has a negative quality, as this not only means that the product is unprofitable to produce, but it sets expectations that products further out beyond the frontier are also unprofitable. The reason niche innovation eventually stops is more novel. It stops for good realizations as well as bad, and when bad, the quality need not be negative. Because niche innovators are competing for existing customers, it is enough that the quality of the new product is so low as to be dominated by neighboring products or so good that it dominates its neighbors.

This setting provides a new perspective on the classic question of how the intensity of competition affects innovation. A consumer's disutility from product differences (the "transportation

cost") is a measure of a market's competitiveness; formally, it measures the substitutability between products and, thus, how much neighboring products compete against each other. For a given set of incumbent products, we show that as competition increases in intensity, the relative appeal of frontier-expanding innovation over niche innovation strictly increases.

This reinforces the idea that competition not only drives more intense innovation, as Arrow argued, but that it also drives more novel innovation. Competition matters, therefore, not only because it drives down costs or allows one firm to steal market share but because it drives the speculative innovation from which true breakthroughs are made, the innovations that expand markets. In examining the relationship between competition and innovation, we thus must look at not only the level of innovation but also the type of innovation that emerges.

REFERENCES

- Acemoglu, Daron.** 1998. "Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality." *Quarterly Journal of Economics* 113 (4): 1055–89.
- 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): 701–28.
- Arrow, Kenneth.** 1962. "Economic Welfare and the Allocation of Resources for Invention." In *The Rate and Direction of Inventive Activity: Economic and Social Factors*, 609–26. Princeton: Princeton University Press.
- Arrow, Kenneth.** 2012. "The Economics of Inventive Activity over Fifty Years." In *The Rate and Direction of Inventive Activity Revisited*, edited by Josh Lerner and Scott Stern, 43–48. Chicago: University of Chicago Press.
- Callander, Steven.** 2011. "Searching and Learning by Trial and Error." *American Economic Review* 101 (6): 2277–2308.
- Callander, Steven, and Niko Matouschek.** 2020. "The Novelty of Innovation: Competition, Disruption, and Antitrust Policy." https://sjc.people.stanford.edu/sites/g/files/sbiybj4051/f/innovation_13october2020.pdf.
- Callander, Steven, Nicolas Lambert, and Niko Matouschek.** 2020. "Congested Innovation in Markets."

- Hotelling, Harold.** 1929. "Stability in Competition." *Economic Journal* 39 (153): 41–57.
- Kim, W. Chan, and Renee Mauborgne.** 2004. "Blue Ocean Strategy." *Harvard Business Review*, October. <https://hbr.org/2004/10/blue-ocean-strategy>.
- Reinganum, Jennifer F.** 1982. "A Dynamic Game of R and D: Patent Protection and Competitive Behavior." *Econometrica* 50 (3): 671–88.
- Romer, Paul M.** 1990. "Endogenous Technological Change." *Journal of Political Economy* 98 (5, Part 2): S71–S102.