

ARTICLE

Innovation on Rugged Technological Landscapes: Implications for Competition Policy

Steven Callander, Nicolas Lambert, and Niko Matouschek

Abstract: Digital and technology markets are often analyzed as if innovation moves along a smooth path. We describe an alternative view: innovation unfolds on rugged technological landscapes, where firms learn from prior successes and failures and choose between expanding the frontier and innovating within existing niches. This distinction helps explain why innovation may stall without exclusion, why product clustering can be ambiguous, and why more innovation is not always the same as more socially valuable innovation. The framework offers competition policymakers a diagnostic tool for interpreting stagnation, entry, differentiation, and the design of remedies, including in debates over ex ante digital regulation.

The present paper offers a condensed version of the article “Innovation and Competition on a Rugged Technological Landscape” forthcoming in the *American Economic Journal: Microeconomics*. For a more in-depth analysis and complete references, readers are referred to the original article.

1. Introduction

Innovation in practice is irregular and uneven. Breakthroughs are rare, failures common, and progress often arrives mixed with reversals. Firms decide not only whether to innovate but also how to innovate. They decide whether to make a leap into the relatively uncontested product space beyond the existing market, or to enter a niche inside the existing market, where opportunities may be better understood but competition is more direct.¹

In recent work, we developed a framework that aims to capture these features of innovation (Callander, Lambert, and Matouschek, 2026a). We conceive of a market

¹ Management scholars refer to these options as *blue ocean* and *red ocean* strategies, respectively (Kim and Mauborgne, 2004).

as built on a *rugged technological landscape*: technological opportunities are unknown, unevenly distributed across the product space, and interconnected. Firms learn about opportunities from what has come before, and this guides both whether they innovate and how they innovate. In a rugged landscape, firms must search for breakthroughs while simultaneously contesting market position against incumbents. The topography of the technological landscape can itself determine the type of innovation that is the more attractive bet for a firm.

The shape that innovation takes matters for competition policy. When innovation is conceived as a smooth, cumulative process, any slowdown looks like a blockage. Remedies naturally then focus on clearing the path. But innovation can stall for many reasons, and not all of them are strategic. The underlying technology and the opportunities it affords shape where innovation happens, what form it takes, and whether it succeeds at all. Strategy and technology are intertwined.

With this framework in hand, we can understand more deeply why innovation thrives and why it stalls, and how innovation and market competition develop jointly, with direct lessons for competition policy. We distinguish between two types of innovation—*frontier* innovation that expands the range of a technology and the market, and *niche* innovation that refines both—and identify differences in their social value. We show how to interpret clustering and differentiation in a market shaped by technological opportunities as well as strategic intent. And we show when and why innovation stalls without exclusion, due to blockages that are technological rather than strategic, identifying policy interventions that can overcome those blockages and foster self-sustaining innovation.

In what follows, we summarize the framework and extract its implications for antitrust practice. The framework shifts the diagnostic questions practitioners should ask. Is a market expanding into new technological territory, or merely filling in existing space? Does clustering reflect anticompetitive conduct, ordinary head-to-head rivalry, or the pull of nearby technological opportunities? Does an innovation slowdown follow a strategic blockage, or a failed attempt to cross a difficult technological valley? We begin by describing what we mean by a rugged technological landscape and how it shapes the logic of market entry, innovation, and competition, then return to these questions as we draw out the lessons for competition policy.

2. What is a Rugged Technological Landscape?

Our model builds on the classic ideas of spatial competition and differentiation (Hotelling, 1929). There is a large set of possible new products that can be developed and introduced to a market. Initially, only a single product is present in the market. The rest of the technological space remains to be explored. These are the opportunities for innovation.

Figure 1 depicts our framework graphically. The space of products is the x-axis (the positive real line). The incumbent product is at 0, and three entrants are located at I_1 , I_2 , and I_3 . Later entrants are free to locate anywhere in the product space. As in

classic spatial models, firms horizontally differentiate along the product space. Consumers have heterogeneous preferences. Each consumer has an ideal product in the space of products and prefers products closer to that ideal.

In addition to horizontal differentiation, firms are vertically differentiated in quality. The quality of a product is initially unknown and given by the technological landscape. This is the black rugged line $v(l)$ in Figure 1. The quality of some products is high, whereas for others it is low, with the full range in between.² 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).

A firm innovates when it introduces a new product into the market. Whether the innovation is high quality and a success, or low quality and a failure, is unknown until the product is tried. The firm does not innovate blindly, however, as it can learn from its predecessors.




Figure 1: Innovation and Competition on a Rugged Technological Landscape

If a product in the market is a success, nearby products are likely to be successes as well, whereas if a product is a failure, nearby products are likely to fail as well, and innovation in that region will stagnate. The expected quality of a new product is given by the dashed line in Figure 1. To form expectations, the firms interpolate between

² The mathematical innovation in our paper is to represent the technological landscape as the realized path of a Brownian motion, with zero drift and positive variance. This is an unusual use of the Brownian motion, although, as we show in our paper, it surprisingly captures the uncertainty of innovation well. The properties we describe in what follows—in particular, the ruggedness—all emerge from the mathematics of the Brownian motion. See Bardhi and Callander (2026) for a survey of research using this approach.

products in the market, and on the market frontier, they extrapolate from the frontier product.

On a rugged technological landscape, the qualities of nearby products are correlated, with the degree of correlation increasing in proximity. This matters for expected quality, and it also matters for risk. Uncertainty is low, and innovation is less risky when it is incremental, but this means that differentiation in the product space is also low, and competitive intensity is high. In contrast, bold and novel innovation far from existing products produces a product that is differentiated from competition, but whether that product succeeds or fails is highly uncertain. This trade-off captures well the distinction between red ocean and blue ocean strategies that is popular in management.

On a rugged technological landscape, firms choose whether to enter the market and, if so, how to innovate. They must choose between innovating on the frontier or in a niche, and in either case, they must decide how novel to be with their innovation, how much to differentiate from what has come before. The uncertainty and impact of an innovation depend on the choice of innovation itself.

This leads to rich dynamics of innovation and competition as documented empirically for new and growing industries (Klepper and Graddy, 1990). Over time, the direction of innovation moves in a cycle: technological setbacks at the frontier push firms into niche innovation, while intensifying competition in niches pushes firms back toward frontier innovation.

If the technological frontier disappoints too much, this cycle can stop. A local low point in the technological landscape is enough to stop frontier innovation. That low point may represent the beginning of a technological valley, with wonderful technological and economic opportunities on the other side. The firms do not know that, however, and their own economic incentives aren't enough to compel them to innovate to find out. The problem is that they will capture only a fraction of the opportunities on the other side of the technological valley that their innovation reveals, with much of the benefit captured by other firms. This creates a context-specific wedge between the private and the social value of innovation.

Even when frontier innovation stops, niche innovation can continue for a long time, but it will ultimately exhaust itself. For innovation to continue indefinitely, a market requires innovation at the frontier, for the technological boundary of the market to be forever expanding. At the frontier, successful innovation begets further innovation. In a niche, in contrast, successful innovation deters follow-on innovation. In fact, the more successful the niche innovation is, the greater the deterrent effect. To thrive and survive, a market needs not only innovation, but it also needs the *right* kind of innovation.

On a rugged technological landscape, market leadership comes apart from technological and quality leadership. A distinction exists between the firm at the leading edge of technology (the right-most product at the boundary of the market),

the firm with the highest quality product, and the firm that leads in profitability. These three forms of leadership can reside within distinct firms. In Figure 1, the firm at I_3 is at the leading edge of the technological frontier, whereas firm I_2 has the highest quality product. Whether one of these firms, or some other firm, is the market leader in profit depends on the firm’s technological advantage as well as the intensity of competition in its region of the product space.

Standard models of spatial competition have always implicitly included a technological landscape –a landscape that is smooth and flat. On a flat technological landscape, firms’ only incentive is to differentiate, to get as far away from their competitors as possible, with each entrant moving further and further onto the frontier. On a rugged landscape, innovation looks very different. Firms turn away from the frontier when it is unpromising, and they do not always seek the most differentiated position even when they could. This shift in perspective –from a smooth landscape to a rugged one– reorders much of what we expect from innovation and competition, with consequences that we draw out in the next section.

Pharmaceutical markets illustrate the distinction. A breakthrough drug may open a new therapeutic pathway and reveal opportunities for follow-on research. A “me-too” or follow-on drug within an existing class may also be valuable, especially if it improves tolerability, dosing, delivery, or competitive pressure. But the two innovations need not have the same social value. Frontier innovation can expand the market and induce further innovation; niche innovation can improve choice within an existing space while also making that space more crowded. For competition policy, the question is therefore not simply whether an industry is producing new products, but what kind of innovative path those products represent.

Artificial intelligence provides a parallel example closer to current digital-policy debates. A firm that pushes the frontier of foundation models –through scale, architecture, training methods, or general capability– may expand the technological boundary and reveal new opportunities for follow-on innovation. A firm that builds a specialized AI tool for legal research, medical coding, software development, or customer support may instead innovate within a narrower space already opened by foundation models. That niche innovation can be highly valuable: it adapts general capabilities to concrete user needs, improves performance in particular domains, and may intensify competition within those domains. But it need not have the same dynamic effect as frontier innovation. The competition policy question is therefore not simply whether the AI sector is producing new products, but whether innovation is expanding the technological frontier or filling in the space behind it.

3. Competition Policy on a Rugged Landscape

A rugged-landscape view of innovation reframes several familiar antitrust questions. The main takeaways are the following.

First, stagnation need not be strategic. Innovation can stall without exclusion. This is a simple but often overlooked point in competition analysis. Even with exclusion and other anticompetitive practices set aside in our model, innovation ebbs and flows and can stagnate purely because of the ruggedness of technological opportunities. The model provides a way to distinguish between technological stagnation and strategic exclusion.

Second, the social value of innovation depends on what type of innovation it is. Frontier innovation begets frontier innovation, whereas niche innovation can deter it. If all innovation is in niches, a market can appear innovative and thriving, yet the gains may be incremental rather than long-lasting and substantial for social welfare. Policymakers should therefore evaluate an industry on whether it is producing new products *and* on how innovative those products are. Our framework offers a way to categorize innovation by its dynamic consequences, not merely by its novelty at the moment of entry.

Third, clustering and differentiation read differently on a rugged landscape. In competition analysis, the location of firms in product space often matters: close clustering may signal head-to-head rivalry, imitation, or, in some settings, concerns about coordinated conduct, while greater differentiation may soften price competition even as it expands variety. Our framework adds another interpretation. When innovating in a niche, a firm may prefer balanced competition in which its nearest neighbors in product space are similar in quality. A firm may choose such a balanced niche even when that niche is narrower, and the expected quality of the new product is lower than the firm could obtain elsewhere. Thus, clustering in product and quality space, even when on lower quality goods, need not by itself reveal anticompetitive conduct, nor does it imply that the resulting pattern of entry is socially optimal. On a rugged landscape, entry and differentiation patterns must be read against the underlying technology.

Fourth, fewer innovations can mean higher welfare. The classic debate of Schumpeter (1942) and Arrow (1962) asks whether competition promotes innovation by disciplining firms or whether it suppresses innovation by reducing the rewards. On a rugged landscape, the answer depends on what kind of innovation competition is induced. In simulations of the model, sharper competitive intensity reduces the number of innovations introduced, yet welfare rises because competition shifts firms toward more socially valuable frontier innovation and consumers extract more value from the products that exist. A related pattern emerges when the technological landscape itself becomes more rugged: fewer innovations occur on average, but the breakthroughs that do occur are more valuable. The lesson is not that fewer innovations are generally better. It is that raw counts of new products are, at best, a partial guide to whether innovation is serving society.

The same interpretative caution applies to who survives. In simulations, frontier innovators account for the majority of all innovations but only a minority of products that remain active when the market matures. Bold frontier innovations, those that are more socially valuable, are more likely to be rendered obsolete by follow-on innovation as they induce their own disruptors to enter the market, succumbing to a Schumpeterian fate. Those that survive necessarily have built successful products, but it need not mean they should be celebrated. Their successful innovations may have been in niches where their success escaped the Schumpeterian fate by deterring rather than fostering further innovation.

Fifth, remedies require a correct diagnosis. To remedy a market failure requires that the failure be diagnosed correctly. If innovation at the frontier stagnates, it need not be the fault of the frontier firm, so clearing the competitive path by breaking up that firm or restraining it in some way may not help innovation restart. The blockage may be the technological landscape itself. Policy tools can help a market cross a technological valley, although doing so requires that existing tools be repurposed and new tools developed.

This diagnosis is easier in the model than in practice. The framework is not a liability screen and does not imply that observed stagnation should be presumed technological rather than strategic. Its value is diagnostic: it identifies a mechanism by which innovation can stall even in the absence of exclusion, and it clarifies what evidence would be needed to distinguish that mechanism from anticompetitive conduct.

This diagnostic point also bears on the new generation of ex ante digital regulation. Regimes such as the Digital Markets Act and analogous initiatives in the United Kingdom, Japan, and Korea begin from the concern that entrenched digital platforms can block competition and innovation. That concern may often be justified. But it also risks importing a smooth-landscape premise: if innovation slows around a dominant platform, app store, or digital gateway, the natural inference is that the platform is the blockage to be removed. A rugged landscape view suggests a further diagnostic step. Before choosing a remedy, policymakers should ask whether the relevant market is stalled because access is being strategically restricted, because the next frontier is technologically difficult to cross, or because innovation has shifted from platform-level innovation into app-level or complementary niche innovation. The point is not to weaken ex ante regulation, but to sharpen the diagnosis on which it rests.

In our paper, we develop a targeted market intervention in which government policy subsidizes innovators when innovation stalls at a technological valley.³ The

³Changing the technological landscape itself is difficult and falls typically within the mandate of research funding agencies such as the NSF, though our work points to natural synergies between the two forms of government involvement.

intervention depends on an important informational condition: the regulator must be able to observe, or at least credibly infer, that frontier innovation has stalled because the market is stuck at a technological valley rather than because of exclusionary conduct, weak demand, or poor firm performance. Under that condition, even small and short-lived subsidies can have outsized welfare effects. Often, subsidies for only one or two firms are all that is necessary to make innovation self-sustaining once again.

In practice, this points less to automatic subsidies than to evidence gathering: agencies and policymakers would need to understand whether firms are abandoning the frontier because rivals are blocking them, because demand is absent, or because the next technological step is privately unattractive but socially valuable.

Targeted subsidies are not the only lever. Familiar tools from antitrust and intellectual property also operate differently in a rugged landscape, and our framework helps clarify how. Permissive treatment of incumbent acquisitions of entrants, for example, can raise the number of innovations but at the cost of their novelty and social value, since the prospect of acquisition dulls an entrant's incentive to innovate boldly and push the frontier (Callander and Matouschek, 2022). Patent protection illustrates a related point, though in a simplified way. In our main paper, we model stronger protection as a longer period before imitation competes away the value of an innovation. In that exercise, stronger protection pushes firms toward bolder and more socially valuable frontier innovation. The result should not be read as a general endorsement of stronger patents in all settings; rather, it shows how a familiar policy lever can change the direction and type of innovation, as well as the amount. In ongoing work (Callander, Lambert, and Matouschek, 2026*b*), we provide the first formalization of patent *width* in a technological space, which similarly compels bolder innovation at the frontier. In each case, the policy lever is familiar, but its operation looks different once we account for the topography of opportunity.

4. Takeaways and Open Questions

The rugged-landscapes perspective does not mean that anticompetitive practices are not salient. Viewing innovation and competition through a rugged-landscape lens complements a purely strategic view. It shows why innovation can stagnate without strategic exclusion, and how strategic incentives can be leveraged—both within the market and through policy interventions—to induce innovation that is more robust and in the form that is most socially beneficial.

For practitioners, the central lesson is diagnostic. On a rugged landscape, familiar policy signals about entry, clustering, differentiation, and innovation slowdowns are less clear than on a smooth one. They must be interpreted against both strategic incentives and the underlying structure of technological opportunity.

Our framework yields specific insights and implications, yet much work remains to be done to put the policy suggestions that follow on more solid ground. This requires careful empirical work to identify where firms innovate, why they stop innovation, and what would be revealed if the government intervened in a market. Combined with further theoretical development, answers to these questions will deliver practical guidance to policymakers and practitioners.

Steven Callander,* Nicolas Lambert, and Niko Matouschek*****

The authors thank Thibault Schrepel for helpful editorial guidance.

* 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

References

- 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.
- Bardhi, Arjada, and Steven Callander. 2026. “Learning in a Correlated World.” *Annual Review of Economics* forthcoming.
- Bloom, Nicholas, Mark Schankerman, and John Van Reenen. 2013. “Identifying technology spillovers and product market rivalry.” *Econometrica* 81(4):1347–1393.
- Callander, Steven, Nicolas Lambert, and Niko Matouschek. 2026a. “Innovation and Competition on a Rugged Technological Landscape.” *American Economic Journal: Microeconomics* forthcoming.
- Callander, Steven, Nicolas Lambert, and Niko Matouschek. 2026b. “Patent Width: Innovation, Competition, and Licensing.” Working paper.
- Callander, Steven, and Niko Matouschek. 2022. “The Novelty of Innovation: Competition, Disruption, and Antitrust Policy.” *Management Science* 68(1):37–51.
- Hotelling, Harold. 1929. “Stability in Competition.” *Economic Journal* 39:41–57.
- Kim, W. Chan and Renee Mauborgne. 2004. *Blue Ocean Strategy*. Harvard Business School Press.
- 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.
- Schumpeter, Joseph A. 1942. *Capitalism, Socialism, and Democracy*. New York: Harper and Brothers.