Defensive Innovation and the Social Return to R&D

Nannan Peng

PhD Candidate, Department of Economics, Rutgers University

June 17, 2026

Extended Abstract

R&D investment is commonly viewed as a key input into productivity growth. Yet in the United States, the relationship between R&D and productivity has weakened over the past two decades: aggregate R&D intensity has continued to rise, while labor productivity and TFP growth have slowed. This paper studies one mechanism that can rationalize this pattern: firms increasingly allocate innovative effort toward defensive innovation, defined as innovation that protects incumbent rents but generates little or no direct productivity improvement. Defensive innovation can have high private value for incumbent firms because it deters rivals and preserves market power, but it may have low or negative social value because it slows follow-on innovation and creative destruction.

The central question of this paper is: how does defensive innovation affect the private and social return to R&D? I answer this question in two steps. First, I provide empirical evidence that patent-based defensive innovation has increased and that it lowers market followers' innovation efficiency. Second, I build a dynamic model in which firms choose between product and defensive innovation. In the model, a decline in the cost of filing patent applications for minor improvements shifts leaders' innovation portfolios toward defensive innovation. This raises total innovation spending but lowers productivity growth because a larger share of innovative effort is devoted to protecting rents rather than improving product quality.

The empirical analysis uses firm-level data from Compustat and PatentView. I propose the share of internal patents as a proxy for defensive innovation. A patent is classified as internal if more than 50 percent of its backward citations are self-citations, where a self-citation refers to a cited patent owned by the same assignee (Akcigit and Kerr, 2018). The economic intuition is that defensive patenting often builds around a firm's existing technologies or products. For example, patent thickets and evergreening strategies frequently involve filing additional patents that cite earlier core patents. These patents may extend protection around existing products without generating proportional improvements in product quality (Frakes and Wasserman, 2025).

I first validate the measure by studying whether leaders' internal patents deter innovation by other firms in the same industry. Specifically, I estimate regressions of followers' patent output on the internal patent share of the technology leader in the same three-digit SIC industry, controlling for year and industry fixed effects, and for follower characteristics such as firm size, capital stock, and accumulated patents. Since patents are count data, I estimate the relationship using both Poisson and negative binomial specifications. The results show a robust negative relationship between leaders' internal-patent share and followers' subsequent patenting. A 10 percentage point increase in the leader's internal-patent share is associated with roughly a 5 percent decline in followers' granted patents. This finding suggests that internal patents capture a strategic component of innovation that reduces rivals' innovation efficiency.

I then document a time-series shift in firms' innovation strategies. Using 2007 as the reference year, I estimate changes in firms' internal-patent share over time while controlling for industry fixed effects and firm-level characteristics. As Figure 1 shows, the share of internal patents is relatively stable before 2007, but rises significantly afterward. By 2020, average internal patenting had increased by about 5 percentage points, corresponding to an increase of more than 100 percent relative to its pre-2007 level. This pattern suggests that firms increasingly shifted their innovation portfolios toward defensive patenting during the same period when productivity growth slowed despite rising R&D intensity.

To interpret these empirical facts, I develop a dynamic model of innovation with leaders and followers based on Akcigit and Ates (2023). Each product market contains two firms that compete in product quality. When one firm has higher quality, it becomes the leader; the other firm is the follower. Leaders choose both product innovation and defensive innovation. Product innovation

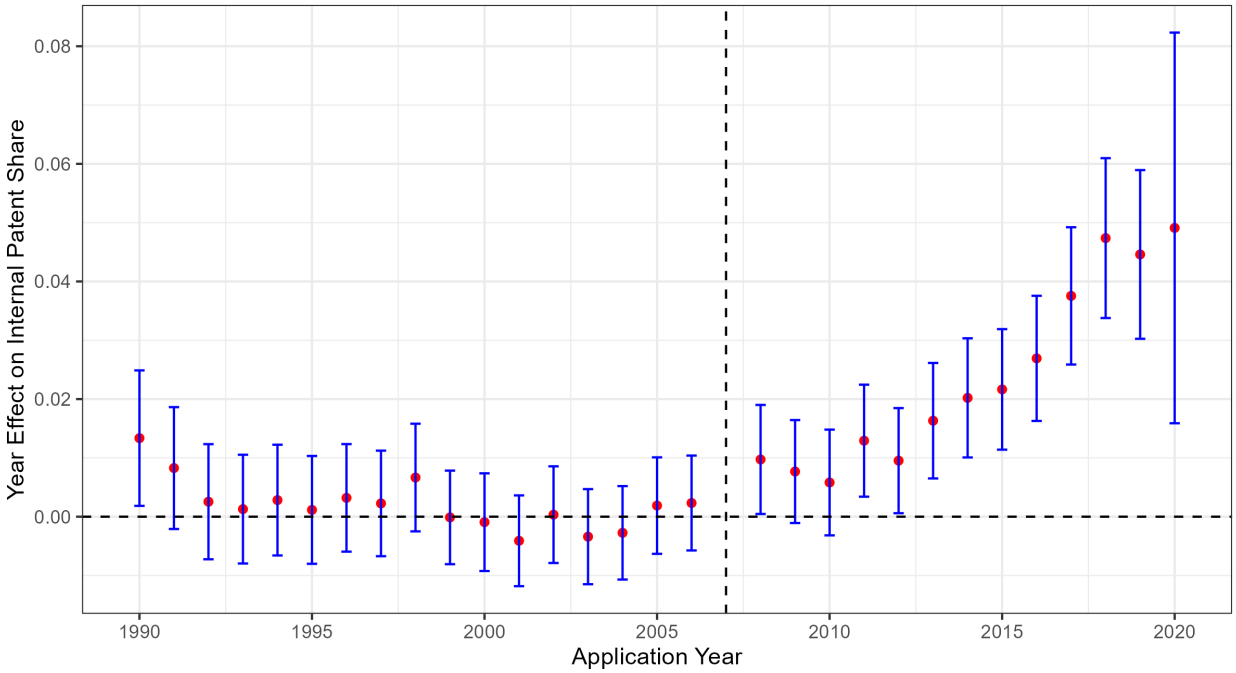


Figure 1: Share of internal patents over time

increases product quality and expands the leader's technological advantage. Defensive innovation increases the leader's stock of defensive patents, which raises the follower's cost of product innovation but does not directly increase product quality. Followers choose only product innovation, but their innovation efficiency decreases as the leader's stock of defensive patents increases. A follower can gradually catch up through incremental innovation or catch up quickly through a drastic innovation shock. When firms are neck-and-neck, they choose product innovation only.

The model highlights a wedge between the private and social returns to R&D. For a leader, defensive innovation is privately valuable because it lowers the probability that the follower catches up. This protects the leader's market share and future profits. From the social perspective, however, defensive innovation is costly because it diverts innovation labor away from productivity-enhancing research and reduces the rate of creative destruction. Therefore, an increase in measured R&D or patenting need not imply a higher contribution to aggregate productivity. The return to R&D depends on the composition of innovation.

A key force in the model is the technological gap between leaders and followers. When leaders are only slightly ahead, product innovation is valuable because it helps them escape close competition. As the gap widens, however, the marginal return to additional product innovation declines. At the

same time, defensive innovation becomes attractive because it helps preserve the leader-follower structure. As a result, leaders with larger technological advantages invest more in defensive innovation. This mechanism is consistent with recent empirical evidence that patents held by large or leading firms are less tightly linked to product introductions and more associated with competitor deterrence (Argente et al., 2020).

I use the model to study the aggregate consequences of a decline in the cost of defensive patenting. This shock captures changes in the legal and institutional environment that make it easier for firms to file patent applications for minor improvements around existing products. After the shock, leaders shift toward more defensive innovation. Total innovation spending increases, but a larger share is used to deter rivals rather than to improve product quality. Followers find it harder to catch up, leader turnover declines, and the firm's distribution shifts toward more persistent leader-follower relationships. Aggregate productivity growth declines because technologically advanced leaders increasingly invest in defensive innovation, which has no direct effect on productivity.

The model also has implications for the wage structure. Innovation workers are used in both productive and defensive innovation, while production workers are used in goods production. When defensive innovation becomes cheaper, the demand for innovation workers remains strong because leaders use them to build defensive patent stocks. At the same time, slower productivity growth lowers production-worker wage growth. Along the transition path, this generates a widening wage premium for innovation workers relative to production workers. Thus, the same mechanism that lowers the social return to R&D can also rationalize the post-2007 rise in the innovation-worker wage premium.

Calibration of the model is ongoing; the quantitative exercise will assess how much of the post-2007 productivity slowdown this mechanism can account for and will provide estimates of the welfare cost of defensive innovation relative to the first-best allocation.

The paper contributes to the measurement of R&D returns by showing that the social return to R&D depends critically on whether innovative effort is directed toward product innovation or defensive innovation. Standard measures of R&D expenditure or patent counts may overstate socially productive innovation when firms use patents strategically to deter rivals. This mechanism also speaks to the patent puzzle: patenting and R&D may remain strongly related to firm-level innovative activity, yet aggregate productivity growth can slow if a rising share of patents protect

incumbent rents rather than generate product-quality improvements (Boldrin and Levine, 2013). The paper also contributes to the literature on declining research productivity, business dynamism, and market power by identifying a specific product-market mechanism through which incumbent firms can protect rents and slow creative destruction. Unlike explanations based only on ideas becoming harder to find (Bloom et al., 2020), or explanations emphasizing that growth has become harder to find conditional on ideas (Fort et al., 2026), this paper emphasizes that firms' innovation choices themselves may shift toward activities with lower social returns.

These results have implications for innovation policy. If part of measured R&D reflects defensive innovation, then policies that broadly subsidize R&D or strengthen patent protection may raise private returns without generating commensurate social gains. A prominent example is the U.S. Section 41 research tax credit, which subsidizes qualifying R&D expenditure across the board without distinguishing between productivity-enhancing and defensive innovation. To the extent that firms direct a portion of this subsidized spending toward defensive patenting, such policies may inadvertently increase the private return to rent-protective activity rather than to socially valuable research. The relevant policy question is therefore not only how much firms invest in R&D, but what type of innovation these investments support. A patent system that permits firms to protect minor modifications around existing products may encourage innovation portfolios that preserve incumbency while reducing follow-on innovation and creative destruction. Understanding the composition of R&D is therefore essential for estimating its true rate of return.

Acknowledgements and Disclosures

I have read the NBER research dissemination disclosure policy. This research did not rely on any external source of funding or research support. I have no relevant or material financial relationships or other potential conflicts of interest to disclose. The views expressed in this paper are my own.

References

Ufuk Akcigit and Sina T Ates. What happened to us business dynamism? *Journal of Political Economy*, 131(8):2059–2124, 2023.

Ufuk Akcigit and William R Kerr. Growth through heterogeneous innovations. *Journal of Political Economy*, 126(4):1374–1443, 2018.

David Argente, Salomé Baslandze, Douglas Hanley, and Sara Moreira. Patents to products: Product innovation and firm dynamics. 2020.

Nicholas Bloom, Charles I Jones, John Van Reenen, and Michael Webb. Are ideas getting harder to find? *American Economic Review*, 110(4):1104–1144, 2020.

Michele Boldrin and David K Levine. The case against patents. *Journal of Economic Perspectives*, 27(1):3–22, 2013.

Teresa C Fort, Nathan Goldschlag, Jack Liang, Peter K Schott, and Nikolas Zolas. Growth is getting harder to find, not ideas. Technical report, National Bureau of Economic Research, 2026.

Michael D Frakes and Melissa F Wasserman. Strategic patenting: Evidence from the biopharmaceutical industry. Technical report, National Bureau of Economic Research, 2025.