New Products

Abhiroop Mukherjee
HKUST

Bruno Pellegrino
Columbia Business School

Alminas Žaldokas
National U. of Singapore

Yiman Ren
UMichigan Ross

Tomas Thornquist
Abu Dhabi Investment Auth.

[Preliminary]

Abstract

Measuring the welfare impact of new product introductions is a major, long-standing problem in economics. In this study, we make progress on this challenge by leveraging the informational efficiency of equity markets and a scalable hedonic consumer demand model. We use stock market reactions to new product announcements over a twenty-year period (2002-2021) to estimate the future profits they generate for their inventor firms. We then use a GHL oligopoly model to measure the change in competitors' profits as well as consumer welfare induced by these new products, ultimately obtaining their welfare contribution in dollars. Our analysis reveals that new products introduced annually by U.S. publicly-traded corporations alone generate substantial welfare gains, averaging 0.15-0.20% of US yearly GDP. Producer surplus accounts for roughly 60% of these gains. This estimate is significantly larger than the corresponding figure for existing products. We show that this is due to the fact that new product creation is highly concentrated among firms that have a high degree of market power.

Correponding Authors:
Abhiroop Mukherjee: amukherjee@ust.hk - HKUST Business School, Clear Water Bay, Kowloon, Hong Kong.
Bruno Pellegrino: bp2713@columbia.edu - Kravis Hall, Office 768, 665 W 130th St, New York, NY 10027 USA.
Alminas Zaldokas: alminas@ust.hk - NUS Business School, 15 Kent Ridge Drive, Singapore 119245.
We thank Andy Atkeson, Jonathan Eaton, Kilian Huber, Chad Jones, Sam Kortum, Hanno Lustig, Song Ma, Elvira Sojli, James Stock, Leo Sveikauskas, Philip Valta, Laura Veldkamp, Lei Xu, Constantine Yannelis and participants at NBER I3, Brookings Institution Productivity Measurement Initiative Authors Conference, European Economic Association Meeting, International Industrial Organization Conference, SGF Conference, Corporate Finance Day, HKUST Interdisciplinary Research Seminar, Baltic Economic Association Meeting, American Economic Association Poster Session, 23rd Taiwan Symposium on Innovation Economics and Entrepreneurship, and Lithuanian Conference on Economic Research for comments, and Benas Zurauskas for excellent research assistance.

1 Introduction

Throughout human history, the composition of economic activity has undergone steady, yet dramatic transformation. The continuous invention and introduction of superior products and services, regularly displacing existing ones, is a fundamental driver of the sustained improvement in living standards that societies have experienced. Schumpeter (1943) famously referred to this process as "creative destruction".

While it is easy to imagine how new product introductions must be contributing significantly (and at all times) to the growth of consumer welfare, measuring this contribution has long been an elusive goal for economists (Bresnahan and Gordon, 1996). This gap in understanding stems from multiple challenges. First, the welfare effect of new products is not captured by real GDP growth, which (by definition) only changes in production volume of existing varieties. Second, quantifying the welfare impact of new products involves significant conceptual and measurement obstacles.

When a new product is invented, it induces a multiplicity of welfare changes for different categories of economic agents: 1) it generates monopolistic profits for its inventor; 2) it generates negative profit spillovers to producers of substitute goods (which lose business to the new product) 3) it generates consumer surplus for customers who purchase those goods; 4) it generates positive spillovers for sellers of existing goods that are strategic complements for the newly-introduced ones. In order to measure the welfare contribution of new products, we must measure all of these effects (the change in total surplus is the sum of all these effects). The magnitude of all these effects in turn depends on the degree of substitution between novel and existing varieties. In other words, to measure the value of new products, we must understand how they interact in the product market (and in the consumer's utility) with existing varieties.

This is well understood in the field of empirical industrial organization (IO). Demand estimation techniques allow IO economists to recover, from observed price, quantity sold, and product characteristics data, the welfare contribution of new products (Hausman, 1996; Petrin, 2002; Nevo, 2003). The main limitation of this approach is that it is only feasible for very few industries where this rich data is available. In other words, it is not scalable for most of the economy. Quantifying the value of new product introductions at scale remains an open research question.

In this paper, we make progress on this question: we develop a methodological framework that allows us to estimate the welfare impact of new product introductions for a very large set of firms, namely the universe of publicly-traded firms. We accomplish this with a three-step approach.

The first step consists of constructing a dataset of new product introductions, and estimating the contribution of those new products on the profit of the firm that introduced them. To do so, we construct a measure of new products, as follows. We obtain the dates of new product announcements from two different sources. The first is S&P CapitalIQ's Key Developments database. As an alternative data source, we use Dow Jones' Factiva database that involve firms traded on U.S. stock exchanges, and that Factiva has filed under New Products/Services category. Using a convolutional neural network trained on a human-generated dataset, we refine this set of article to a much smaller one that captures with high likelihood new product announcements. We then estimate abnormal returns on the mentioned firm's stock on the day of the media articles, and consider only those firms with a positive stock price reactions. This ensures that our procedure captures new product releases that are important enough to move firms' stock prices.

The second step in valuing the welfare contribution of new products is to rely on financial markets to identify important innovation. The stock price impact of a media article in a short time-window surrounding its publication reflects the market's estimation of all current and future profits from the new product mentioned in that article. In an efficient stock market, this is the correct estimate of the value of the product to the firm's shareholders. To the extent that expected profits for the innovating firm reflects the aggregate

willingness-to-pay for that good or service by consumers, this value of the product to the firm, in turn, reflects the value created by that particular innovation. Such approach is similar to Kogan, Papanikolaou, Seru and Stoffman (2017), who value patents using stock market reactions to patent grants. The market-based measure of value is not only flexible – in that the same approach can be used to value any type of good, service, or activity – it also helps avoid issues associated with researchers trying to measure value of new business models, say, a ChatGPT 4.0 or a Meta virtual reality headset.

The third consists in computing the product market spillovers of new product introductions, including the impact on consumer surplus. In order to accomplish that, we need a demand system that can cover all U.S. publicly-traded firms. We use the GHL oligopoly model of Pellegrino (forth.), which is ideally suited to this purpose. The model allows us to estimate the impact of the product introduction not only on the firm itself, but also on its competitors and on consumers.

Using this approach we find that new products generate substantial welfare gains, averaging 0.15-0.20% of GDP annually. We find that most of the welfare gain from new products accrues to the innovating firms, rather than to the consumer. Further investigation reveals that the key reason behind such a pattern is that new products are introduced disproportionately by firms with substantial market power, who then extract most of the surplus from this process in the form of increased profits.

Looking at the cross-section, we find that product innovation generates significant spillovers to producers of substitute goods as well as strategic complements. Our methodology allows us to quantify these spillovers for each individual firm pair. We find significant heterogeneity in the sign and size of these spillovers across firms.

We also find interesting cyclical patterns. The welfare contribution of new products declined significantly during the 2008-2009 financial crisis, possibly as a consequence of the fact that tight financial conditions limited firms' ability to invest in product innovation. The size of the product market spillovers, on the other hand, spiked around the COVID pandemic, reflecting the fact that product innovation was concentrated among an unusual set of firms, which likely were active in more competitive markets.

Our analysis also reveals two key patterns in how innovation benefits are distributed. First, product innovation is disproportionately concentrated among firms with substantial market power, explaining why producers capture a notably larger share of the welfare gains compared to what is typically observed for existing products. Second, the economic impact of new products is highly concentrated, with the top 10% of innovations consistently generating over 85% of both profits and welfare gains. While this concentration has gradually declined from about 95% in the early 2000s to around 86-88% by 2021, the fact that such a small fraction of new products accounts for the vast majority of economic benefits underscores the disproportionate impact of "breakthrough" product innovations.

Previous Literature. Our paper connects to three important streams of literature, belonging respectively to Industrial Organization (IO), Finance, and Macroeconomics.

First, we build on a pre-existing literature in IO that seeks to understand new products on consumers. Bresnahan and Gordon (1996) dedicated an NBER volume to the challenges of measuring and valuing new goods; in that volume, Hausman (1996) and Greenstein (1996) used demand estimation tools to study the introduction of (respectively) new varieties of ready-to-eat cereals and computers. In subsequent work, Petrin (2002) quantified the welfare benefits of the introduction of minivans by using modern empirical IO methods. These studies present detailed analyses of isolated, specific products and markets. While leveraging a similar conceptual toolbox, our study differs significantly in that it seeks (more ambitiously) to carry out a similar exercise at *scale*. Namely, it seeks to value dozens of thousands of products that span virtually all industries.

The second key literature that our paper connects to consists of stock market event studies of new product introductions and other innovation-related events. Our paper is not the first to link equity market valuations

to new products. Eddy and Saunders (1980) studied the impact of new product introductions on monthly stock returns using a sample of 66 firms, followed by Wittink et al. (1982) on computer and office machines business. Chaney, Devinney and Winer (1991) study new product introductions over 1975-1984 and find an average stock price reaction of 0.75% over a 3-day window. Austin (1993) uses an event-study approach to value biotech innovations, while Sood and Tellis (2009) study five industries in electrical products. Chen et al. (2005) show a negative stock price effect on rivals. While our use of a stock market event study connects us to this literature, this study seeks to answer a very different research question: we are not interested per se in establishing whether new products generate abnormal equity returns (as this can already be inferred from existing studies). Instead, in this paper we want to push the idea that we can utilize stock market reactions to value individual new products and (by leveraging a structural oligopoly model) even establish their welfare impact.

In this respect, our paper borrows significantly, on the methodological side, from Kogan, Papanikolaou, Seru and Stoffman (2017, henceforth KPSS), who used stock market reactions to value patents. It deviates from it, however, in important respects: 1) the focus on new products, as opposed to patents and 2) the expanded focus on not just the profits generated by the new products, but on the product market and consumer spillovers. Pakes (1985) also provides an early contribution examining the relation between patents and the stock market rate of return. Aside from KPSS, stock market reactions have recently been used more broadly in economics, for instance by Amiti, Gomez, Kong and Weinstein (2024): they study the welfare effects of trade wars by examining market responses to tariff announcements.

The third important literature that our paper connects to lies in the domain of macroeconomics. New products have long been of interest to macroeconomists due to their importance for the measurement of inflation and the theory of ideal price indices. In an early seminal study, Feenstra (1994) developed a methodology to adjust price indices for the entry and exit of new varieties, with an application to import prices and income elasticities. Broda and Weinstein (2006, 2010) and Aghion, Bergeaud, Boppart, Klenow and Li (2019) expand Feenstra's framework in order to, respectively, the bias to overall inflation and GDP growth that is induced by unmodeled product introductions and discontinuations. In related work Garcia-Macia, Hsieh and Klenow (2019) decompose TFP growth into various components (own innovation, creative destruction, and the introduction of new varieties) using a model of growth with creative destruction. The common denominator of these studies is that they restrict their focus on the *aggregate* impact of new products on inflation and TFP, and that they infer their role using indirect measurement through the lens of a model.

Our paper, while related, departs from this literature in three important respects. Firstly, like IO studies, it is data-driven, as opposed to model driven (our analysis begins by analyzing data on newly-announced products). Secondly, it studies the impact of new products at a granular level (as opposed to aggregate). The third difference lies the research question: we seek to create new firm-level indicators of profits, consumer surplus and product market spillovers generated by new products, in contrast with the previous studies that seek to correct or decompose headline macroeconomic indicators such as TFP or inflation.

This study also contributes to an emerging literature that leverages large GHL oligopoly models (Pellegrino, forth.; Ederer and Pellegrino, 2021; Eeckhout and Veldkamp, 2023; Hopenhayn and Okumura, 2025). Finally, it relates to the seminal study by Gort and Klepper (1982) on the diffusion of 46 new products, as well as certain other studies that focus the econometric identification of spillovers (see Bloom, Schankerman and Van Reenen, 2013 and Huber, 2023 among others).

2 Methodology: New Product Announcements

2.1 Conceptual Challenges and Definitions

Our empirical analysis requires a clear definition of what constitutes a "new product" and when we consider a product to be "new." This seemingly straightforward question reveals profound conceptual challenges that merit careful consideration.

Defining "New Products". We adopt an inclusive definition of "new products" that encompasses not only entirely novel product introductions but also significant updates, improvements, and upgrades to existing products. This broad definition reflects the reality that innovation often occurs incrementally, with firms enhancing and refining their existing product offerings over time.

Timing Challenges. A fundamental conceptual challenge is determining *when* a product becomes "new." This can occur at multiple times along a continuum, including: 1) when first conceived in the mind of inventors; 2) when actively developed within the firm; 3) when publicly introduced to the market; 4) when available to consumers; 5) or when widely utilized by consumers.

At each point in this continuum, the expected profit stream from the same product may differ substantially. As information becomes available, technical feasibility is established, consumer interest is gauged, and competitive responses materialize, the expected value of the innovation continuously evolves. This creates a moving target where no single, stable measure of a product's expected profit can be considered definitive.

For different product types, different points in this continuum might offer the most accurate value assessment. For highly anticipated annual product iterations, such as Apple's iPhone, much of the value may be priced into company stocks well before announcement, with the announcement effect capturing only the surprise relative to expectations. For truly novel products, the announcement might capture a more substantial portion of the innovation's value.

Our Measurement Approach. Recognizing that there is no universal solution to this challenge, we make a methodological choice to focus on the product announcement stage as our measurement point. We define a product as "new" at the moment the company publicly announces it to the market. This approach offers several advantages: it provides a consistent, observable event across diverse product types; it typically represents the first point at which substantial concrete information becomes publicly available; it allows us to leverage efficient market reactions to estimate expected value; and it is methodologically tractable for systematic large-scale analysis.

While this choice is imperfect—and counterexamples like highly anticipated product refreshes are readily apparent—it represents a pragmatic compromise that works reasonably well across a wide range of product innovations. The stock market reaction at announcement captures the market's revision of expected future profits based on newly available information about the product.

Modeling Profit Dynamics. Beyond defining when a product is "new," another key challenge is how to model the profit contribution of new products over time. Products may take time to build market share, face changing competitive conditions, or require ongoing investments in marketing and improvement. Rather than attempt to model these complex dynamics, we make the simplifying assumption that each new product generates a constant stream of profits until it is discontinued. This profit stream represents the product's contribution in an average year over its lifecycle.

While this constant-profit assumption abstracts from the rich dynamics of product lifecycles, it provides a tractable way to translate stock market reactions into welfare estimates. When investors value a new product announcement, they are effectively pricing the present value of all future profits generated by that product. By combining this market-based estimate with our empirical analysis of product survival rates,

we can recover the implied average annual profit contribution without having to specify the precise path of profits over time.

This approach allows us to measure the aggregate welfare contribution of new products while remaining agnostic about the specific timing of profits within a product's lifespan. The key identifying assumption is not that profits are literally constant, but rather that stock market participants can rationally price the present value of the entire future profit stream at the time of product announcement.

2.2 CapitalIQ Key Developments

Our source of new product announcements data is S&P Capital IQ's Key Developments database. This database tracks significant corporate events, including new product introductions, across publicly traded companies. We specifically focus on event type 41 (Product-Related Announcements), which captures various forms of product-related news including product launches, unveilings, and major enhancements.

To ensure we capture genuine new product introductions rather than routine corporate announcements, we apply several filters to the raw data. First, we restrict our attention to headlines containing specific action words that signal new product introductions: "unveil", "reveal", "launch", "introduce", "release", "present", "issue", "upgrade", "enhance", "improve", "extend", or "new." We explicitly exclude announcements related to natural resource exploration (containing words such as "drill", "explore", or "excavate") to avoid confounding new product introductions with operational updates in extractive industries.

A crucial step in our methodology is to isolate the effect of new product announcements from other corporate events that might affect stock prices. To this end, we remove announcements that coincide with major corporate events in a three-day window centered on the announcement date. We exclude announcements that occur within one day of earnings announcements (identified using Compustat's earnings announcement dates), merger and acquisition announcements (sourced from SDC Platinum), and dividend declarations or stock splits (obtained from Capital IQ's corporate actions database). This cleaning procedure helps ensure that our measure captures the market's reaction to new product introductions rather than other coincident corporate events that might affect stock prices.

2.3 Estimating the Profit Impact of New Products

In our next critical step, we leverage the efficient market hypothesis to obtain a forward-looking measure of the private value of new products for every firm/year. We estimate abnormal returns on the mentioned firm's stock surrounding the day of the article's release. Specifically, we calculate cumulative abnormal returns on the day of the product announcement. Expected returns are calculated based on the Fama-French 3-factor model. We excluded days if the firm in the article announced earnings or an M&A transaction on that day (including one day before and one day after for both events) as these major events might confound our estimates.

This method is similar to the one taken by Kogan et al. (2017) (KPSS) to value patents, and it allows us to capture the market's immediate reaction to the new product announcement, which should reflect the expected present value of future profits generated by the new product. It is important to note that while our approach is inspired by KPSS, there are a few key differences. We focus on new products, which are the final outputs of innovation that directly reach consumers. This allows us to capture innovation that might not be patentable or types of innovation that firms might choose not to patent for strategic reasons.

Linking announcements to their stock market value ensures that we have a market-based measure of product value, which is both forward-looking in nature (given that the stock market's reaction to an announcement

FIGURE 1: PRODUCT ANNOUNCEMENT - EVENT STUDY

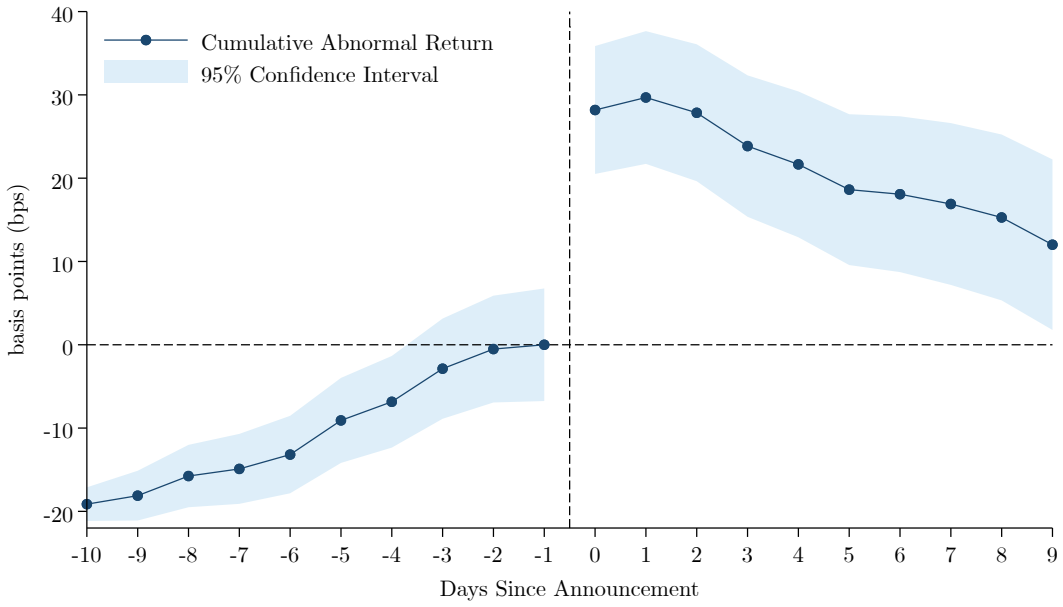


FIGURE NOTES: This figure shows cumulative abnormal returns (CARs) around new product announcements, estimated using the Fama-French three-factor model. The solid line represents the average CAR, while the dashed lines represent the 95% confidence intervals. The event window spans from 10 days before to 9 days after the announcement ($t = 0$), with CAR normalized to zero at $t = -1$. CARs are expressed in basis points (one basis point equals 0.01%). The sample includes all new product announcements from both Capital IQ and Factiva that were identified by our methodology and that did not coincide with other major corporate events (earnings announcements, M&A, or dividend/split announcements) in a three-day window around the announcement.

accounts for all future profits or losses from it), and avoids issues associated with the researcher figuring out what type of product actually adds value – be it an app or an appliance.

Figure 1 presents the event study analysis of cumulative abnormal returns (CARs) around new product announcements. The analysis is based on the Fama-French three-factor model and examines a window of 20 trading days, centered on the announcement date (from $t = -10$ to $t = +9$). We normalize the CAR to zero at $t = -1$ to focus on the announcement effect.

The results reveal a sharp and economically significant market reaction to new product announcements. On the announcement day ($t = 0$), we observe an immediate jump in CARs of approximately 29 basis points. This positive reaction continues into the following day (t=1), reaching a peak of about 30 basis points. The magnitude of this two-day response suggests that new product announcements convey substantial information about firms' future prospects.

The positive market response persists in the post-announcement period, though it gradually declines. By ten days after the announcement (t+9), CARs remain positive at around 13 basis points, indicating that the market's initial reaction is not reversed in the short term. The 95% confidence intervals (shown by the dashed lines) indicate that the positive CARs are statistically significant during the announcement window and remain so throughout most of the post-announcement period.

Looking at the pre-announcement period ($t - 10$ to $t - 1$), we observe a slight downward trend in CARs, from about -19 basis points ten days before the announcement to our normalization point of zero at $t = -1$.

This pattern could suggest some anticipation effects or reflect broader market dynamics leading up to new product announcements.

2.4 From Abnormal Returns to Expected Yearly Profits

The event study analysis yields an estimate of the change in a firm's market capitalization as it introduces a new product. However, what we need for our measurement exercise is a measure of yearly profits that are generated by the corresponding product. While the two are naturally connected by the fact that a firm's market value is the discounted present value of all its future profits, to infer the latter from the first we need to make an assumption.

We make the conservative assumption that all new product introductions induce a parallel shift in the term structure of future dollar profits. Under this assumption, the contribution of new product z to next period profits (v_z) can be estimated as:

$$v_z \;=\; r_{it(z)}\,\delta_z\,\mathcal{M}_{it(z)-1} \tag{2.1}$$

where δ_z is the percent increase in the market capitalization of company i induced by product introduction z, $\mathcal{M}_{it(z)-1}$ is the previous period market capitalization, and $r_{it(z)}$ is a measure of the discount rate implied by the firm's market capitalization.

Now, it may be tempting to use the measured abnormal return - which we call $\mathcal{A}_{it(z)}$ - to proxy for δ_z. However, the abnormal return $\mathcal{A}_{it(z)}$ is bound to contain, along with information about the market value change induced by the introduction of product z, a significant amount of measurement error $\varepsilon_{it(z)}$. The standard and natural assumption (as suggested by KPSS) is that this error is additive:

$$\mathcal{A}_{it(z)} \;=\; \delta_z + \varepsilon_{it(z)} \tag{2.2}$$

One undesirable consequence of measurement error is that, if we were to use $\mathcal{A}_{it(z)}$ as a proxy for δ_z, we would infer that the profit impact of many new products is negative. By accounting for the presence of this measurement error, we can obtain a superior estimate of the economic contribution of the new products. Following KPSS, we assume that the signal follows a half-normal distribution with variance parameter τ_δ^2

$$\delta_z \;\sim\; \mathcal{N}^{+}\left(\tau_\delta^2\right) \qquad \varepsilon_{it} \;\sim\; \mathcal{N}\left(0, \tau_\varepsilon^2\right) \tag{2.3}$$

defining ρ to be the signal-to-noise ratio

$$\rho \;=\; \frac{\tau_\delta^2}{\tau_\delta^2 + \tau_\varepsilon^2} \tag{2.4}$$

we obtain the mean estimate of δ_z, conditional on the observed abnormal return $\mathcal{A}_{it(z)}$:

$$\mathbb{E}\left(\delta_z \middle|\, \mathcal{A}_{it(z)}\right) \;=\; \rho\,\mathcal{A}_{it(z)} + \sqrt{\rho}\cdot\tau_\varepsilon \cdot \frac{\varphi\left(-\sqrt{\rho}\cdot\delta_z\cdot\tau_\varepsilon^{-1}\right)}{1-\Phi\left(-\sqrt{\rho}\cdot\delta_z\cdot\tau_\varepsilon^{-1}\right)} \tag{2.5}$$

by plugging this expectation inside equation (2.1) we obtain our (non-negative) estimate of v_z.

Using this data, we create two measures of new product introductions at the firm-year level. Our first measure is a count of the number of new products launched by a firm in a given year. This is defined as the number of distinct times we uncover, using the approach described above, news media mentions of new products associated with the firm that lead to positive abnormal returns in the stock market. The positive abnormal return condition ensures that we are capturing products that are important enough for the firm mentioned in our articles to have a meaningful impact on its value (and consequently rules out cases where the article mentions new products introduced by rivals, or related firms; these are unlikely to affect a firm's

return in the same manner).

For our second measure, instead of counting new products, we aggregate cumulative abnormal returns (conditional on them being positive, as before) coming from all new product news for a given firm in a given year. Note that the annual horizon of aggregation is to capture the fact that many consecutive news articles within a time span might refer to the same new product launch by a firm, with each articles revealing a bit more information on the product (and hence, each article being accompanied by positive stock price reactions). Aggregating over a longer time span, such as two years instead of one to capture even slower diffusion of information about new products in the stock market does not change our results materially.

Note also that our positive abnormal return screen can mitigate the concern that the trends that we observe might be related to changes in the media coverage of firms over time. Still, in all our regression specifications we also use time fixed effects to account for such possibilities (assuming that any media coverage changes are not systematically different across sectors over time).[1]

We aggregate over an annual horizon to capture the fact that multiple news articles within a time span might refer to the same new product launch, with each article revealing more information. To translate these abnormal returns into dollar values, we multiply the cumulative abnormal return by the firm's market capitalization. This gives us an estimate of the market value added for each new product announcement.

3 Demand Model, Spillovers and Welfare

3.1 Generalized Hedonic-Linear (GHL) Demand

The next step in our methodology is to estimate, starting from the measured profit contributions of new products, the corresponding product market spillovers and welfare contributions. In order to accomplish that, we need a demand system and a model of product market competition that can encompass all US publicly-traded firms. The standard and natural tool in this case is the Generalized Hedonic-Linear (GHL) Oligopoly model of Pellegrino (forth.), which we review briefly below.

Time is discrete and indexed by t. There are n firms in the economy, indexed by $i \in \{1, 2, ..., n\}$, that produce differentiated products and compete oligopolistically. Each firm produces only one product, i.e., firm i only produces product i. The demand these firms face is hedonic, meaning that consumers evaluate each product as a bundle of characteristics.

There are two types of characteristics: m common characteristics indexed by $k \in \{1, 2, ..., m\}$ and n idiosyncratic characteristics indexed by $i \in \{1, 2, ..., n\}$. Each unit of product i yields one unit of the corresponding idiosyncratic characteristic i a unit-length vector of $\mathbf{a}_i \in \mathbb{R}^n$ of common characteristics, formally:

$$\mathbf{a}_{it} = \begin{bmatrix} a_{1it} & a_{2it} & \dots & a_{nit} \end{bmatrix}' \tag{3.1}$$

$$\text{such that} \qquad \sum_{k=1}^{m} a_{kit}^2 = 1 \quad \forall\, i \in \{1, 2, ..., n\} \tag{3.2}$$

Because each unit of good i provides exactly one unit of its corresponding idiosyncratic characteristic, we can just use q_i to denote the units of idiosyncratic characteristic i consumed by representative consumer.

We assume that the representative consumer's utility function $\mathcal{U}$ is linear quadratic in both common characteristics $\mathbf{x}_t$ and idiosyncratic characteristics $\mathbf{q}_t$, and that it is linear (decreasing) in the labor supply,

[1] Another alternative would be to scale the number of new product announcements by the extent of media coverage about the firm. However, this method would assume that the ratio of new product coverage to total media coverage is constant across firms, which is unlikely to be the case.

L_t:

$$\mathcal{U}(\mathbf{x}_t, \mathbf{q}_t, L_t) \quad \stackrel{\text{def}}{=} \quad \sum_{i=1}^{n} b_{it} q_{it} - \frac{\alpha}{2} \cdot \sum_{k=1}^{m} x_{kt}^2 - \frac{1-\alpha}{2} \sum_{i=1}^{n} q_{it}^2 + L_t \tag{3.3}$$

where b_i is the intercept demand shifter for product i and α is the weight placed by the consumer on the common characteristics. The representative consumer earns income from labor and the profits of all firms, and chooses an optimal consumption bundle $\mathbf{q}_t$ subject to a budget constraint. Firm i faces a marginal cost of production equal to c_i and chooses output q_i to maximize profits. Pellegrino (forth.) shows that the following linear demand system solves the consumer problem:

$$\mathbf{p}_t \;=\; \mathbf{b}_t - (\mathbf{I} + \mathbf{\Sigma}_t)\, \mathbf{q}_t \tag{3.4}$$

where $\mathbf{p}_t$, $\mathbf{q}_t$ and $\mathbf{b}_t$ are vectors of prices, quantities, and demand intercepts (respectively). $\mathbf{I}$ is the identity matrix, and

$$\mathbf{\Sigma}_t \quad \stackrel{\text{def}}{=} \quad \frac{\partial \mathbf{p}_t}{\partial \mathbf{q}_t} \quad = \quad \alpha\,(\mathbf{A}_t' \mathbf{A}_t - \mathbf{I}) \tag{3.5}$$

as explained in Pellegrino (forth.) the (i, j) entry of matrix $\mathbf{A}_t' \mathbf{A}_t$, equal to $\mathbf{a}_{it}' \mathbf{a}_{jt}$, is called the "cosine similarity" between firm i and firm j, and is a measure of product substitution. The Cournot-Nash equilibrium vector of profits $\boldsymbol{\pi}_t$ is given by:

$$\sqrt{\boldsymbol{\pi}_t} \quad = \quad \mathbf{\Omega}_t\,(\mathbf{b}_t - \mathbf{c}_t) \tag{3.6}$$

$$\text{where} \qquad \mathbf{\Omega}_t \quad \stackrel{\text{def}}{=} \quad (2\mathbf{I} + \mathbf{\Sigma}_t)^{-1} \tag{3.7}$$

$\mathbf{b}$ and $\mathbf{c}$ are, respectively, the demand and supply function intercepts. Hence, $(b_i - c_i)$ is the marginal surplus of the very first unit produced by firm i. The equilibrium Consumer Surplus is given by:

$$S_t \;=\; \frac{1}{2} \cdot (\mathbf{b}_t - \mathbf{c}_t)'\, \mathbf{\Lambda}_t\, (\mathbf{b}_t - \mathbf{c}_t) \tag{3.8}$$

where

$$\mathbf{\Lambda}_t \;\stackrel{\text{def}}{=}\; \mathbf{\Omega}_t\; (\mathbf{I} + \mathbf{\Sigma}_t)\; \mathbf{\Omega}_t \tag{3.9}$$

3.2 Measuring Product Market Spillovers

We now consider how product introductions can change equilibrium profits and consumer surplus. We assume (without loss of generality) that product introductions do not affect the marginal cost of production c_{jt}.[2] Then (by total differentiation) the change in profit from $t-1$ to t can be written as:

$$\pi_{jt} \;\approx\; 2\sqrt{\pi_{jt}} \left(\sum_{j=1}^{n} \omega_{ijt}\, \Delta b_{jt} \;+\; \sum_{j=1}^{n} b_{jt}\, \Delta \omega_{ijt} \right) \tag{3.10}$$

where Δ indicates the 1-period difference. The first term in parentheses reflects the parallel shifts in residual demand generated by the the introduction of new products. The second reflects the effect of changes in the degree of similarity between products.

We next make an important assumption that the second term in parentheses is small:

$$\sum_{j=1}^{n} b_{jt}\, \Delta\, \omega_{ijt} \;\approx\; 0 \tag{3.11}$$

[2] A 1 dollar change in b_{jt} has the same effect on firm j's profits as a 1 dollar decrease in c_{jt}. This difference is not material for our measurement exercise.

This assumption requires that the cosine similarity $\mathbf{A}'\mathbf{A}$ is approximately constant over the time interval considered for the measurement (one year) and/or that its increments are broadly uncorrelated with $\mathbf{b}$. What this assumption means, intuitively, is that we can model new product introductions, over short time periods, as vertical innovations – which affect the demand intercept $\mathbf{b}$ but not $\mathbf{A}'\mathbf{A}$. In other words, over small (one-year) time horizons, new products are upgraded versions of existing products.

Under this assumption, the expression for the change in profit of firm j simplifies to:

$$v_z \;\approx\; 2\sqrt{\pi_{it(z)}} \sum_{j=1}^{n} \omega_{ijt(z)}\, \Delta b_{jt(z)} \tag{3.12}$$

Although (by assumption) new product introductions will be reflected in changes in $\mathbf{b}$, and we do recover changes both ω_{ij} and Δb_{jt} by taking the model to the data, it would be imprudent to attribute all the time variation in $\mathbf{b}$ to new product introductions. It is important that we allow $\mathbf{b}$ to be influenced by other factors such as random shocks, which could originate for instance by changes in consumer taste.

Let us index new products introduced by firm i with the generic subscript z. Let also $\mathcal{Z}_{it}$ be the set of new products introduced by firm i, between $t-1$ and t. We model changes in b_i over time, as the sum of the effect of the introduction of new products (β_{zt}) and of a random shock ξ_{jt} which captures all other economic forces affecting willingness to pay b_{jt} (the underlying process is irrelevant for our measurement):

$$\Delta b_{jt} \;=\; \sum_{z \in \mathcal{Z}_{it}} \beta_{zt} + \xi_{jt} \tag{3.13}$$

where β_{zt} is the contribution of new product z, introduced at time t. We can then define v_{iz} the effect of introduction of product z on firm j as:

$$v_{jz} \quad \stackrel{\text{def}}{=} \quad \frac{\partial \pi_{jt}}{\partial b_{jt}} \cdot \beta_{zt} \quad = \quad 2\sqrt{\pi_{jt}} \cdot \omega_{jjt} \cdot \beta_{zt} \tag{3.14}$$

v_{jz} is what we actually estimate using stock returns in the previous step of our procedure. Specifically, our analysis of abnormal return results in a dataset of estimates of profit impacts (v_{jz}) for all covered products.

More generally, we can define v_{iz}, the effect of the introduction of z on a generic company i (not necessarily equal to j, who introduced the product):

$$v_{iz} \quad \stackrel{\text{def}}{=} \quad \frac{\partial \pi_{it}}{\partial b_{jt}} \cdot \beta_{zt} \quad = \quad 2\sqrt{\pi_{it}} \cdot \omega_{ijt} \cdot \beta_{zt} \tag{3.15}$$

Then, combining equations (3.14) and (3.15) we can finally write the spillover to firm j in terms of the private rent v_{jz} and other observables ($\boldsymbol{\pi}$ and $\boldsymbol{\Omega}$):

$$v_{iz} \quad = \quad \sqrt{\frac{\pi_{it}}{\pi_{jt}}} \cdot \frac{\omega_{ijt}}{\omega_{jjt}} \cdot v_{jz} \tag{3.16}$$

Finally, we consider the effect on consumer surplus S_t, which we indicate with the capital letter V:

$$V_z \quad \stackrel{\text{def}}{=} \quad \frac{\sum_{i=1}^{n} (b_{it} - c_{it})\, \lambda_{ijt}}{2\sqrt{\pi_{jt}} \cdot \omega_{jjt}} \cdot v_{jz} \tag{3.17}$$

The sum of all these effects (own profits, competitors' profits as well as consumer surplus) is the change in total surplus resulting from a new product introduction. If we then aggregate across all new products introduced in a given year we obtain a measure of the dollar change in welfare from year to year due to the

FIGURE 2: VALIDATION USING TRADEMARKS AND RAVENPACK PRODUCTS

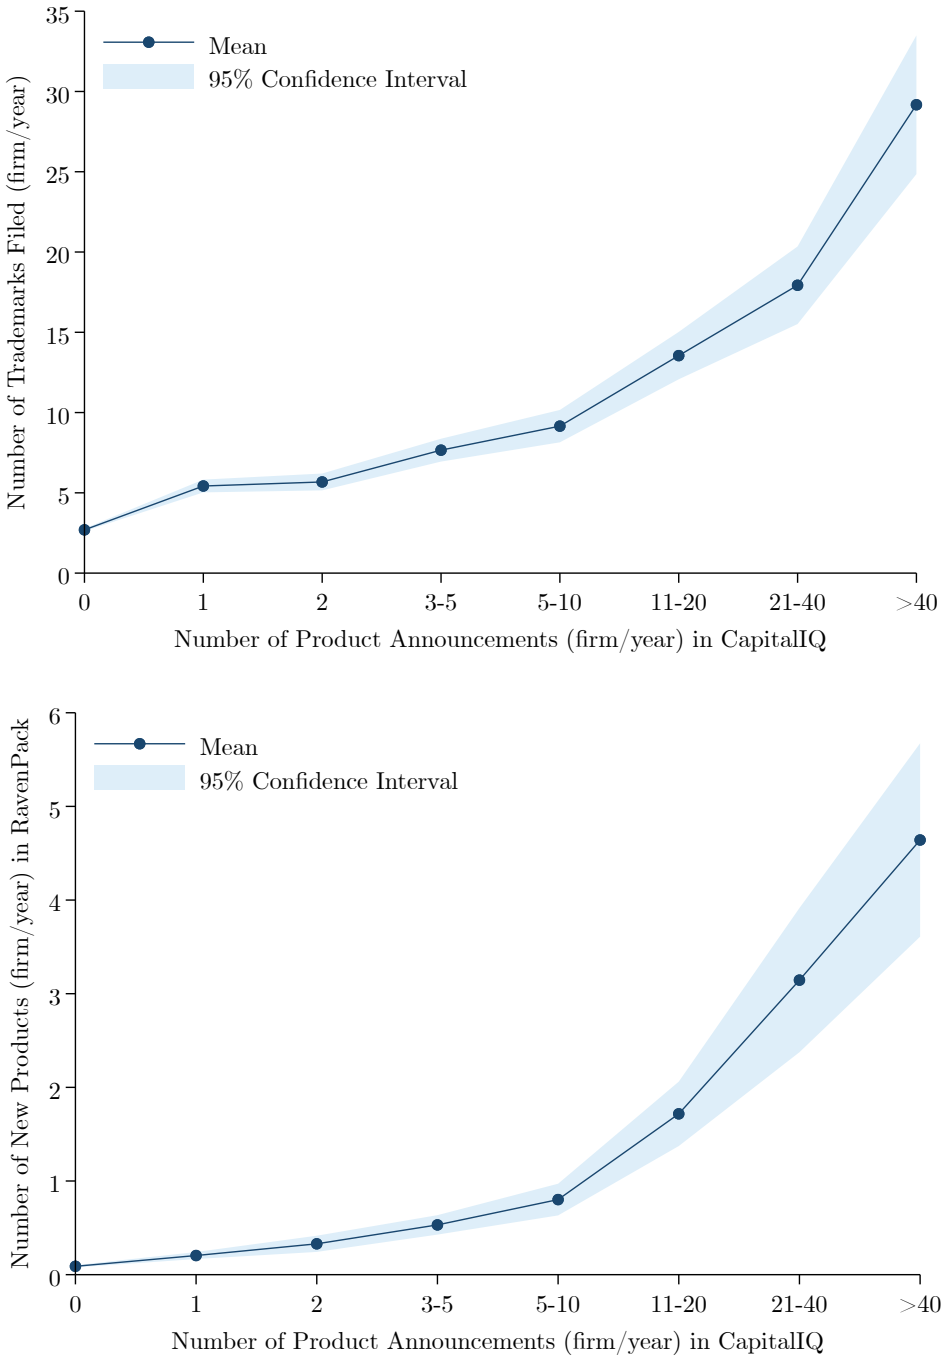


FIGURE NOTES: The figures above shows the relationship between product announcements in Capital IQ and trademarks (upper figure) and product announcements in Capital IQ and product introductions identified from RavenPack dats (lower figure) over 2002-2019. The horizontal axis shows bins for the number of product announcements per firm-year in Capital IQ. The vertical axis shows the average number of new products identified by RavenPack for firms in each Capital IQ announcement bin. A product is considered new in RavenPack when it is first associated with a company in news coverage. The blue dots represent means while the light blue bands show 95% confidence intervals.

introduction of new products.

3.3 Data and Implementation

There are two crucial inputs required to estimate GHL demand. The first is revenue and cost data for publicly-traded companies from Compustat. Based on this data Pellegrino (forth.) shows how to recover the prices $\mathbf{p}$, the quantities $\mathbf{q}$, and the marginal costs $\mathbf{c}$.

The second crucial input is the dataset of Hoberg and Phillips (2016): it provides a time-varying empirical estimate of $\mathbf{A}_t'\mathbf{A}_t$, and is constructed from textual descriptions of products contained in the mandatory annual 10-K filings of the public corporations. Armed with this estimate, we can recover the matrices $\mathbf{\Sigma}_t$, $\mathbf{\Omega}_t$, $\mathbf{\Lambda}_t$ as well as the vector of demand intercepts $\mathbf{b}$ (see Pellegrino for full details).

4 Validation and Relation to Other Measures of Innovation

4.1 Validation using RavenPack and Trademark Filings

Unlike patents or R&D expenditures, there are no comprehensive, well-established databases tracking the dynamics of product introductions across companies over time. This creates a validation challenge: how can we ensure that our new dataset of product announcements from Capital IQ accurately captures meaningful product innovation?

We first look at the data on trademarks that that are filed with US Patent and Trademark Office (USPTO) when a firm launches a new product line or service.[3] The upper figure in Figure 2 presents a validation exercise, in which we plot the average number of trademarks filed by firms with different frequencies of Capital IQ product announcements, over the period 2002-2019. The relationship between the two measures is strongly monotonic: firms with more product announcements in Capital IQ tend to have more trademarks.

As our second and a more direct way to validate our measures, we construct an independent product-firm panel using RavenPack Analytics, a dataset of news articles from over 19,000 sources. RavenPack employs natural language processing algorithms to tag news articles with various entity identifiers, including "company" and "product" tags. A key advantage of using RavenPack for validation is that its product tags provide independent verification of whether something is truly a product – if something appears in RavenPack with a product tag, we can be highly confident it is indeed a product. However, RavenPack has two important limitations. First, it identifies a product as "new" in the year when it is first tagged together with a company in a news article, which may occur well after the actual product announcement. Second, RavenPack's coverage is incomplete – many genuine products are never tagged in its system. While these limitations explain why RavenPack identifies far fewer products than Capital IQ (16,732 versus 97,096 over our sample period), RavenPack's independent verification of product status makes it valuable for validating our main dataset.

Similarly to trademarks, the bottom figure in Figure 2 plots the average number of product introductions identified by RavenPack for firms with different frequencies of Capital IQ product announcements, over the period 2002-2019. The relationship between the two measures is strongly monotonic: firms with more product announcements in Capital IQ tend to have more new products identified by RavenPack. For instance, firms with no announcements in Capital IQ average close to zero new products in RavenPack, while firms with more than forty announcements per year in Capital IQ average about 5 new products annually according to RavenPack.

[3] We thank the authors of Hsu et al. (2021b,a)for sharing this data.

Table 1: New Product Introductions vs. Patenting

		Receives a Patent		Total
		No	Yes	
Announces New Products	No	3,681 (39.0%)	722 (7.7%)	4,403 (46.7%)
	Yes	2,182 (23.1%)	2,848 (30.2%)	5,030 (53.3%)
Total		5,863 (62.2%)	3,570 (37.9%)	9,433 (100.0%)

Table Notes: This table presents the joint distribution of patenting activity and new product announcements for U.S. publicly-traded firms over the period 2002-2021. The sample consists of 9,433 unique firms. "Announces New Products" is a dummy variable equal to one if a firm has at least one new product announcement in the CapitalIQ database. "Receives a Patent" is a dummy variable equal to one if a firm was granted at least one patent during the sample period, based on the USPTO patent data assembled by Kogan, Papanikolaou, Seru and Stoffman (2017). Numbers in parentheses show the percentage of firms in each cell relative to the total sample.

This strong correlation between two independently constructed measures provides additional validation for our Capital IQ-based sample of new product announcements. The fact that RavenPack identifies fewer products than Capital IQ announces is expected given RavenPack's limitations in timing and coverage. However, the strong relationship between the two measures, combined with RavenPack's reliable product identification, suggests that Capital IQ is indeed capturing genuine product introductions.

4.2 New Product Introductions vs. Patents and R&D

A natural question is how our database of new product announcements relates to patent data, which has been extensively used in prior literature to measure innovation. In particular, Kogan, Papanikolaou, Seru and Stoffman (2017) develop an influential approach to measuring innovation by combining patent grants with stock market reactions. Like our methodology, KPSS leverage financial markets to estimate the private value of innovation. However, while patents capture the creation of new technologies, our measure focuses on the introduction of new products - the final stage of the innovation process where firms actually bring novel offerings to market. To understand how these two measures relate to each other, we examine their joint distribution.

Table 1 presents the joint distribution of patenting activity and new product introductions at the firm level over the period 2002-2021. The results reveal a striking pattern: while patenting and new product introductions are positively correlated (correlation coefficient of 41.4%), there is substantial divergence between these two measures of innovation.

Perhaps the most notable finding is that 23.1% of firms in our sample introduce new products without receiving any patents during this period. This substantial group of non-patenting innovators highlights a key limitation of using patents alone to measure innovation: many firms engage in product innovation without seeking patent protection for their innovations. This could reflect various factors, including strategic decisions to protect intellectual property through secrecy rather than patents, innovations that are not patentable (such as new business models or service innovations), or incremental improvements that build on existing

Table 2: Correlation with other Measures of Innovation

Panel A: R&D

	All Sectors (1)	Industrials (2)	Consumer (3)	Healthcare (4)	Financials (5)	ITC (6)	Others (7)
R&D	0.163***	0.084***	0.073*	0.257***	0.134**	0.163***	0.039*
	(0.013)	(0.029)	(0.039)	(0.019)	(0.068)	(0.026)	(0.023)
Firm F.E.	Y	Y	Y	Y	Y	Y	Y
Year F.E.	Y	Y	Y	Y	Y	Y	Y
R^2	0.588	0.508	0.493	0.454	0.608	0.681	0.437
N	67,741	7,798	11,840	18,386	1,983	19,737	7,978

Panel B: Patents

	All Sectors (1)	Industrials (2)	Consumer (3)	Healthcare (4)	Financials (5)	ITC (6)	Others (7)
Patents	0.012**	0.017	0.023	0.036***	0.073***	0.018**	0.006
	(0.005)	(0.012)	(0.015)	(0.012)	(0.026)	(0.009)	(0.012)
Firm F.E.	Y	Y	Y	Y	Y	Y	Y
Year F.E.	Y	Y	Y	Y	Y	Y	Y
R^2	0.611	0.528	0.471	0.464	0.573	0.668	0.361
N	129,931	14,727	18,715	17,763	21,496	24,038	21,290

Figure Notes: The table is constructed from regressions estimated in a panel setting at a firm level. In all specifications, the outcome is the number of new products over the period 2002-2021. Panel A presents results where we correlate new product measure with R&D expenditures as reported in Compustat. Panel B presents results where we correlate new product measure with the number of new eventually-granted patents that the firm has filed in the year as reported in Kogan, Papanikolaou, Seru and Stoffman (2017). We split the observations by GICS sectors: Column (1) pools all sectors, Column (2) reports Industrials (GICS=20), Column (3) reports Consumer (GICS=25&30), Column (4) reports Healthcare (GICS=35), Column (5) reports Financials (GICS=40), Column (6) reports Information Technology and Communication Services (GICS=45&50), and Column (7) reports Other Sectors. Standard errors are clustered at the firm level.

technologies.

At the same time, 7.7% of firms receive patents but do not announce any new products. These firms might be engaging in more basic research or developing technologies that they license to others rather than commercializing directly. The largest group in our sample (39.0%) consists of firms that neither patent nor introduce new products during this period.

Notably, 30.2% of firms both receive patents and introduce new products, indicating that these two innovation metrics often complement each other. However, the fact that over half of the innovating firms in our sample (23.1% out of 53.3%) would be missed by looking at patents alone underscores the importance of our new product-based measure in capturing a more complete picture of corporate innovation activity.

These findings suggest that patent counts, while informative, may significantly underestimate innovation activity by missing firms that innovate through channels that do not require or warrant patent protection. Our new product announcement measure thus provides important complementary information about innovation activity, particularly in sectors where patenting may be less relevant or strategic secrecy more valuable.

FIGURE 3: AGGREGATE WELFARE CONTRIBUTION OF NEW PRODUCTS (2001-2021)

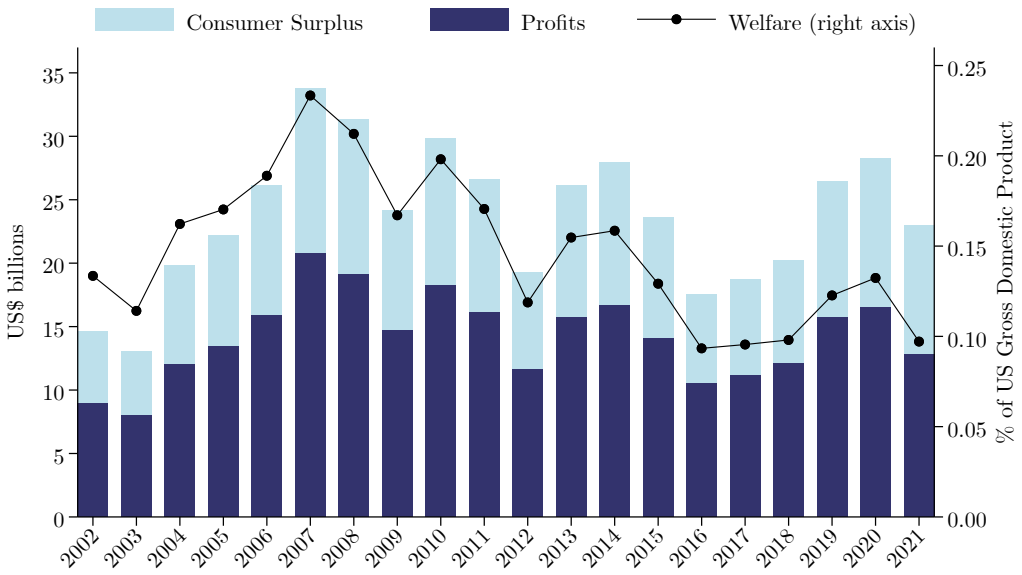


FIGURE NOTES: This figure shows the aggregate welfare contribution of new product introductions by publicly-traded U.S. firms from 2002 to 2021. The stacked bars represent the composition of welfare gains, decomposed into producer surplus (profits, shown in dark blue) and consumer surplus (shown in light blue), measured in billions of U.S. dollars on the left axis. The black line, scaled on the right axis, shows the total welfare contribution as a percentage of U.S. GDP.

4.3 Correlation with Other Measures of Innovation

We now examine how our new products measure correlates with the other constructs of innovation. We consider innovation inputs such as R&D investments and intermediate outputs such as patents. We gather R&D expenditure data from Compustat and patent data from USPTO with a firm-level match provided by Kogan et al. (2017).

In Table 2 we estimate such contemporaneous correlations between our annual new new product measure and R&D as well as patents in a firm-level panel regression with firm and year fixed effects. Panel A reports the correlations for R&D and Panel B reports the correlations for patents. Looking at the columns (1) across both panels, we find that both R&D and patents are correlated with new product introductions. The latter is consistent with the evidence in Balasubramanian and Sivadasan (2011) who link patent stock to the product information in Census data.

We further look at how these contemporaneous correlations vary by industry sector. In Table 2, Panel A, columns (2)-(7), we see that for R&D a contemporaneous correlation holds for all sectors. In Panel B, we look at patents and we see contemporaneous correlation for Financials, Healthcare and, Information Technologies, suggesting that for these sectors traditional measures of innovation could be capturing product innovation trends quite well, but this might not hold universally across all sectors.

Figure 4: Aggregate Welfare Decomposition (2001-2022)

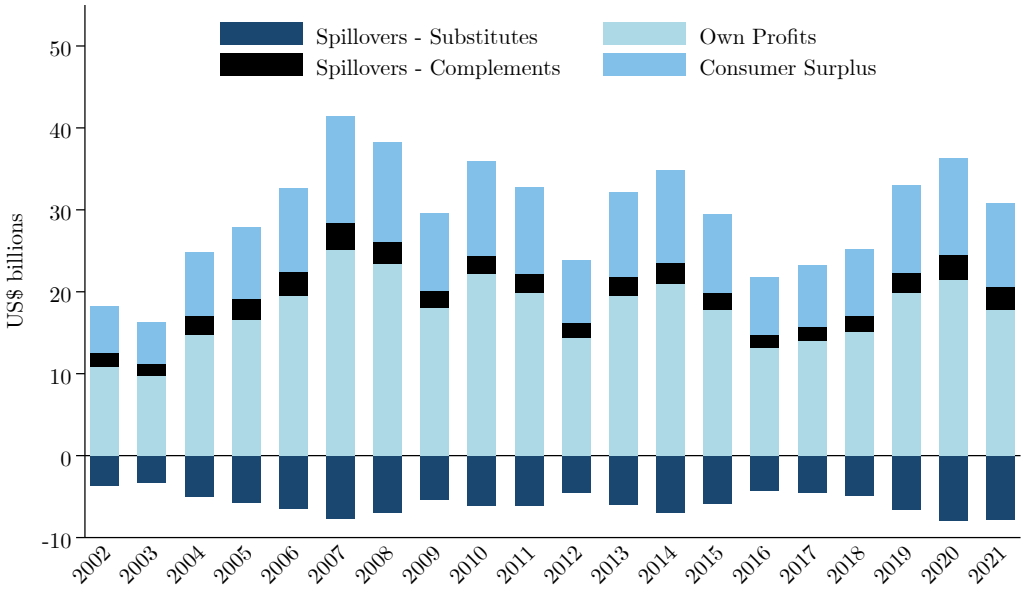


Figure Notes: this figure decomposes the aggregate welfare contribution of new product introductions by U.S. publicly-traded firms into four components over the period 2002-2021. The stacked bars show: (1) consumer surplus (mid blue), representing the direct benefit to consumers; (2) complementary profits (black), capturing positive spillover effects on firms producing complementary products; (3) substitution effects (dark blue), measuring negative spillovers on competing firms; and (4) own profits (light blue), representing direct profits accruing to innovating firms. All values are expressed in billions of current U.S. dollars.

5 Results

5.1 Welfare Contribution of New Products

Figure 3 presents our estimates of the aggregate welfare contribution of new products introduced by publicly-traded U.S. firms from 2002 to 2021. The welfare effects are decomposed into producer surplus (profits) and consumer surplus, and are presented both in absolute terms (billions of U.S. dollars) and as a percentage of U.S. GDP. Welfare estimates are calculated using our three-step methodology: first identifying valuable new products through media coverage and stock market reactions, then estimating their profit impact through event studies, and finally computing welfare effects using the GHL demand system. Both producer and consumer surplus are expressed in current prices.

Our analysis reveals that new product introductions generate substantial welfare gains, with total benefits ranging from 0.10% to 0.24% of GDP over the sample period. The welfare contribution peaked in 2007, when new products generated approximately $21.6 billion in profits and $13.6 billion in consumer surplus, totaling about 0.24% of GDP. Throughout our sample period, the distribution of welfare gains remains relatively stable, with producer surplus consistently accounting for roughly 60% of the total benefits, while consumer surplus comprises the remaining 40%.

The temporal pattern of welfare contribution exhibits significant cyclicality and sensitivity to macroeconomic conditions. This is particularly evident during the 2008-2009 financial crisis, when we observe a sharp decline in both components of welfare. Producer surplus fell from $19.8 billion in 2008 to $15.3 billion in 2009, with consumer surplus following a similar pattern, dropping from $12.6 billion to $9.8 billion. The

FIGURE 5: HETEROGENEITY OF SPILLOVERS TO COMPETITORS AND CONSUMER (2021)

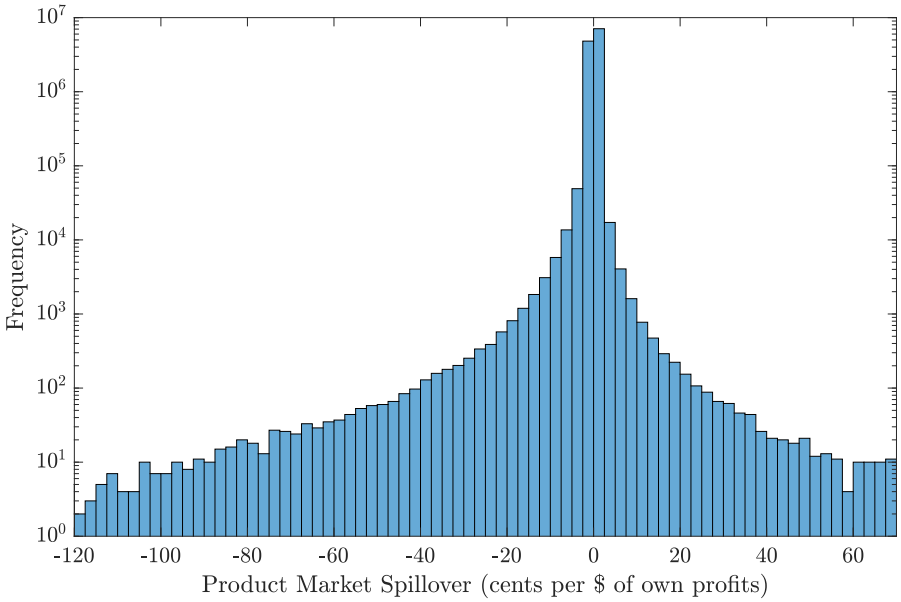


FIGURE NOTES: this figure shows the distribution of product market spillovers generated by new product introductions in 2021, measured in cents per dollar of the innovating firm's own profits, while the vertical axis shows the frequency on a logarithmic scale. Spillovers are estimated using our GHL demand system and capture how each new product introduction affects the profits of other firms in the market. The distribution is based on all new product introductions by publicly-traded U.S. firms in 2021 that generated significant stock market reactions.

market subsequently recovered, but the post-2015 period shows generally lower levels of welfare contribution compared to the pre-2015 era.

By the end of our sample period in 2021, new products generated approximately $13.5 billion in profits and $10.7 billion in consumer surplus, amounting to about 0.10% of GDP. This represents a notable decline from the peaks observed in the mid-2000s, suggesting potential changes in the innovation landscape or in the ability of firms to capture value from their new product introductions. These findings quantify the substantial economic impact of new product introductions and highlight the importance of innovation for both producer and consumer welfare.

5.2 Spillovers

Figure 4 decomposes the aggregate welfare effects into four distinct components: consumer surplus, complementary profits (positive spillovers to other firms), substitution effects (negative spillovers to competitors), and own profits of innovating firms. This decomposition reveals the complex interplay between value creation and value capture in product markets.

Our analysis shows that the direct profits captured by innovating firms form the largest component of welfare gains, averaging around $18-20 billion annually during our sample period. However, these gains are partially offset by negative spillovers to competitors through substitution effects, which typically amount to $5-7 billion per year. This substitution effect reflects the business-stealing aspect of new product introductions, where innovative firms capture market share from their competitors.

Interestingly, we also find substantial positive spillovers to complementary firms, averaging $2-3 billion annually throughout most of the sample period. These complementary effects arise when new products enhance the value of related products or services offered by other firms. The magnitude of these positive spillovers is particularly noteworthy in the later years of our sample, with a marked increase during 2020-2021 to around $7 billion annually, possibly reflecting the growing importance of product ecosystem effects in the modern economy.

Consumer surplus remains a significant and stable component of total welfare, typically ranging between $8-12 billion annually. This relatively stable share of consumer surplus suggests that firms have maintained consistent pricing power over their new products throughout our sample period, despite changes in market structure and competitive dynamics.

Figure 5 presents the distribution of product market spillovers across firms in 2021, measured in cents per dollar of innovating firms' own profits. The histogram reveals substantial heterogeneity in how new product introductions affect other market participants. The distribution is roughly bell-shaped but slightly asymmetric, with a longer left tail, indicating that negative spillovers can be more extreme than positive ones.

The majority of spillover effects cluster around zero, with the highest frequency of observations falling in the range of -10 to +10 cents per dollar of innovating firm profits. However, the distribution extends significantly in both directions, ranging from approximately -120 to +70 cents per dollar. This wide range suggests that while most new products have relatively modest effects on other firms, some innovations can generate substantial positive or negative externalities.

5.3 Concentration

Having established the substantial aggregate welfare contribution of new products, we now examine how these gains are distributed across the spectrum of product innovations. Specifically, we investigate the degree of concentration in both profits and welfare effects by analyzing what share of total gains is generated by the most successful new products. This analysis provides important insights into whether the benefits of product innovation are broadly distributed or concentrated among a small subset of particularly impactful new offerings.

Our analysis reveals striking levels of concentration in both the profits and welfare generated by new products. Figure 6 shows that throughout our sample period, the top 10% of new products account for an overwhelming share of both total profits and welfare gains, though this concentration has declined somewhat over time.

At the beginning of our sample in 2002, the top 10% of new products generated approximately 95% of total profits and 94% of total welfare gains. This extraordinary concentration persisted at similar levels through the late 2000s, with the share hovering around 94-95% for both measures. However, starting around 2011, we observe a gradual but consistent decline in concentration.

The downward trend accelerated after 2015, with the share of profits captured by the top 10% of new products falling from about 92% to 88% by 2021. The concentration of welfare effects shows an even more pronounced decline, dropping to approximately 86% by 2021. This represents a cumulative decrease of about 7-8 percentage points in concentration over the sample period.

Notably, throughout most of the period, the concentration of profits remains slightly higher than the concentration of welfare effects, typically by 1-2 percentage points. This small but persistent gap suggests that the most successful new products tend to generate somewhat more concentrated benefits for their producers than for the broader economy.

Figure 6: Value of New Products: Concentration

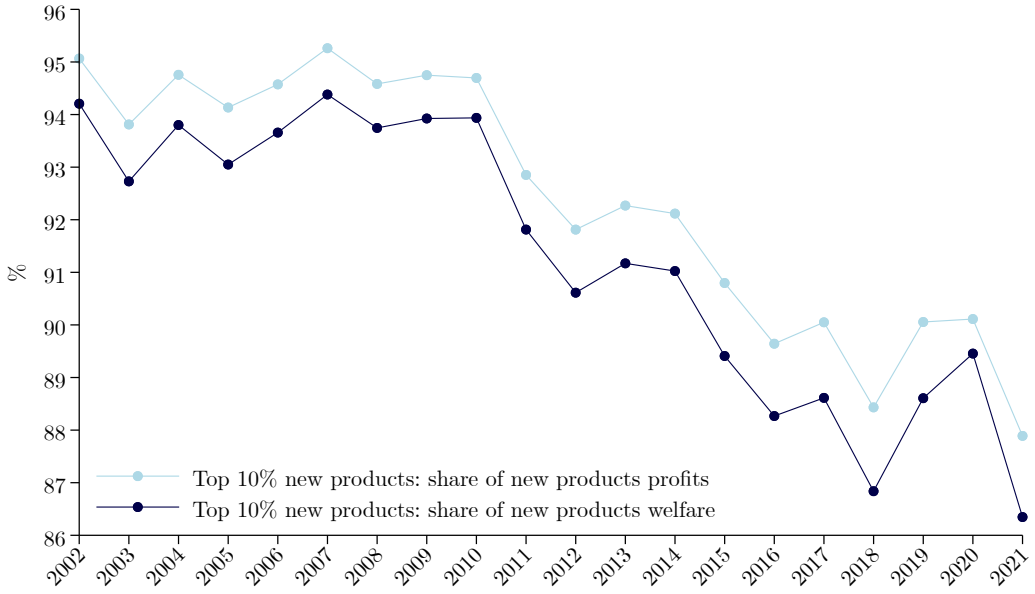


Figure Notes: This figure shows the share of total profits and welfare generated by the top 10% of new product introductions by U.S. publicly-traded firms over the period 2002-2021. The light blue line represents the share of total new product profits accounted for by the top decile of products (ranked by profit contribution), while the dark blue line shows the share of total welfare gains generated by the top decile of products (ranked by welfare contribution). Both measures are expressed as percentages.

The declining concentration over time could reflect several factors, including increased competition in product markets, more widespread capacity for innovation across firms, or changes in the nature of new products themselves. However, despite this modest democratization of innovation benefits, the degree of concentration remains remarkably high, with the top 10% of new products continuing to account for the vast majority of both profits and welfare gains.

5.4 New Products Introductions and Market Power

An important question that emerges from our welfare analysis is why the consumer surplus generated by new product introductions appears relatively small compared to producer surplus, with a ratio of roughly 40:60. This is particularly striking when compared to the much larger consumer surplus share generated by existing products in steady state, which prior research has shown to be around 80% of total surplus. What explains this discrepancy?

The answer lies in understanding where in the product market network new product creation occurs. To measure firms' market power, we employ the measure of product market centrality (χ_i) developed by Pellegrino (forth.) – a metric that captures a firm's competitive position in the network of product market rivalries. Firms with low centrality face less competition from substitutes and thus behave more like monopolists, while highly central firms face intense competition and behave more competitively. Figure 7 provides key insight into this question by showing the distribution of product market centrality across firms, both unweighted and weighted by the number of new product announcements.

The contrast between these two distributions is striking. The unweighted distribution (light blue area)

Figure 7: Distribution of Product Market Centrality

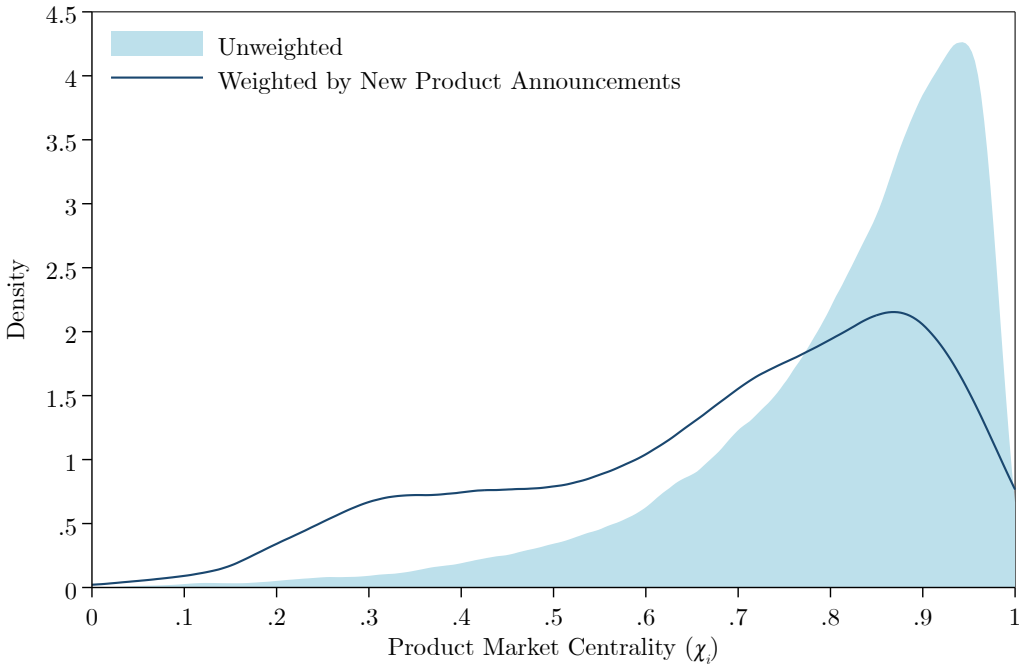


Figure Notes: This figure shows two distributions of product market centrality (χ_i) for U.S. publicly-traded firms in 2021. Product market centrality measures a firm's competitiveness based on its position in the network of product market rivalries, ranging from 0 (monopolist-like behavior) to 1 (perfect competition). The light blue area shows the unweighted distribution across all firms, while the dark line shows the distribution weighted by the number of new product announcements per firm.

shows that most firms in the economy have relatively high centrality ($\chi_i > 0.7$), indicating they operate in competitive product markets. However, when we weight firms by their new product announcements (dark line), we observe a marked shift in the distribution toward lower centrality values, with significantly more mass in the range of $\chi_i = 0.2 - 0.6$. This indicates that firms with greater market power – those facing relatively little competition from substitute products – are more active in introducing new products.

This pattern helps explain why new products generate relatively low consumer surplus compared to producer surplus. Firms with low centrality can behave more like monopolists when pricing their innovations, allowing them to capture a larger share of the surplus created. The fact that these less central firms are more active in introducing new products naturally leads to a lower share of surplus accruing to consumers. In contrast, the higher consumer surplus share for existing products reflects that most firms in the steady state operate in more competitive markets where their ability to extract surplus is more limited.

6 Robustness and Limitations

6.1 Timing of Information Release and Investor Learning

Our measurement approach relies on two key timing assumptions that warrant further discussion. First, we assume that the market's reaction to product announcements fully captures the value created by new products. However, information about new products may be revealed gradually over time rather than being

fully incorporated in stock prices at the time of announcement. For instance, the full market potential of a new product may only become apparent after its launch, or related innovations (like patents) may be announced separately. If information diffusion is indeed gradual, our event-study methodology would only capture a portion of the total value created by new products. This would make our welfare estimates conservative, as we would be systematically understating the profit impact of product innovations.

A second, potentially offsetting concern is that product announcements may cause investors to update their beliefs not only about the announced product itself, but also about the firm's future innovative capacity. In other words, the stock price reaction we measure might reflect both the value of the new product and an upward revision in expectations about the rate of future innovation by the company. This could lead us to overstate the welfare contribution of the specific products we identify.

However, evidence on how investors process information about firms' innovation capabilities suggests our estimates are unlikely to be significantly inflated by belief updating about future innovation. Cohen, Diether and Malloy (2013) document that investors are slow to recognize the implications of firms' past R&D success, consistently undervaluing the R&D investments of firms with strong track records. This systematic under-reaction to information about innovation ability suggests that any updating of beliefs about future innovation potential at the time of product announcements is likely to be incomplete rather than excessive. While stock prices may partially reflect updated expectations about future innovation, the documented tendency of markets to under-react to signals about R&D productivity makes it unlikely that our welfare estimates are substantially overstated due to belief updating effects.

6.2 Sample Selection

Another potential concern with our methodology is that we only observe products that firms choose to release, which could create sample selection bias. Specifically, firms likely receive signals about the potential profitability of products under development and choose to release only those products with sufficiently positive signals. This means our sample of observed products is positively selected on both fundamental value and noise in these signals.

However, this selection process should not bias our welfare estimates if market participants are rational. When forming demand for a firm's stock following a product announcement, investors understand that firms only release products when their private signals exceed a certain threshold. They will therefore adjust their valuations to account for this selection, incorporating both the positive fundamental selection and the positive noise selection. In efficient markets, the stock price reaction will reflect the expected value of the product conditional on the firm choosing to release it. While this conditional expectation is higher than the unconditional expectation across all potential products (including those never released), it is precisely the relevant quantity for measuring the welfare contribution of products that actually reach the market.

In other words, even though released products represent a selected sample, rational market pricing ensures that our stock market-based measurement approach captures the true expected value of these products. The sample selection inherent in firms' product release decisions does not lead us to systematically overstate or understate the welfare gains from new products.

6.3 Physical Complementarities

As noted by Bresnahan and Gordon (1996), a key challenge in measuring the welfare contribution of new products is accounting for complementarities between goods. The value created by a new product often depends on its degree of complementarity with other products and services.

One important advantage of GHL demand is that, thanks to its flexibility, it is able to capture *strategic* complementarities: this is the reason why we measure positive spillovers. That is, when a new product is introduced, the price of another product may increase as a consequence of the fact that said product and the newly-introduced product share some common competitors.

There is, however, another type of complementarity that is not captured by our framework, which is *physical* complementarity: this is a situation where one additional unit supplied of a certain product directly increases the marginal utility of one unit of another product. This occurs for a product pair (i, j) when σ_{ij} is negative. Because the cosine similarity data of Hoberg and Phillips (2016) is positive by construction, one limitation of GHL is that (as of today) it only accommodates physical complementarity in theory, not in the data. The absence of physical complementarities in the GHL framework likely leads to understating the total welfare gains from new products, particularly for innovations that create value primarily through enhancing the usefulness of existing goods. Future work extending the model to incorporate physical complementarities would provide a more complete picture of the welfare effects of product innovation.

7 Conclusions

This paper introduces a novel methodology to measure the welfare contribution of new product introductions at scale. Our approach combines two key innovations: first, we develop a systematic way to identify valuable new products by analyzing media coverage and stock market reactions across the universe of publicly traded U.S. firms; second, we use a scalable demand system to translate these market reactions into welfare estimates, accounting for both direct effects and market spillovers.

Our analysis yields several important findings. First, we document that new products generate substantial welfare gains, averaging around 0.15-0.20% of GDP annually during our sample period. These gains are split between producer surplus (roughly 60%) and consumer surplus (roughly 40%), suggesting that firms are able to capture a significant portion of the value they create through innovation. This relatively low share of consumer surplus, compared to what is typically observed for existing products, reflects a crucial feature of the innovation landscape: new product creation is disproportionately concentrated among firms with low product market centrality – those facing limited competition from substitutes and thus able to exercise significant market power.

Second, our decomposition of welfare effects reveals complex patterns of value creation and redistribution in product markets. While innovating firms capture substantial profits from their new products, we find significant negative spillovers to their competitors, reflecting the "creative destruction" aspect of innovation. We also document substantial positive spillovers to firms producing complementary products, particularly in recent years, highlighting the growing importance of product ecosystem effects in the modern economy. We find considerable heterogeneity in how new products affect different market participants. While most product introductions have modest spillover effects on other firms, some innovations generate large positive or negative externalities, with effects ranging from -120% to +70% of the innovating firm's own profits. This heterogeneity underscores the importance of accounting for market spillovers when assessing the welfare implications of innovation.

Third, we document two distinct patterns in how innovation benefits are distributed. On one hand, unlike for existing varieties, the majority of the welfare gains accrue to the producers in the form of monopoly profits. We show that this is likely a consequence of the fact that product innovation tends to be carried out by firms with significant market power. On the other hand, these gains are remarkably concentrated among a small number of breakthrough innovations, with the top 10% of new products consistently generating over 85% of both profits and total welfare benefits. While this concentration has gradually declined from about 95% in the early 2000s to around 86-88% by 2021, it remains strikingly high.

Finally, we show that the welfare contribution of new products appears to be tied to macroeconomic environment. We observe a notable decline during the 2008-2009 recession (when financial conditions were tight), coupled with a large spike in the size of the product market spillovers during the COVID pandemic, when product innovation was concentrated more than usual in competitive sectors.

Our findings have important implications for both research and policy. For researchers, we provide a new framework for studying product innovation that can be applied consistently across industries and over time. For policymakers, our results highlight the substantial welfare gains from product innovation while also drawing attention to its redistributive effects. Particularly noteworthy is our finding that the innovation process is dominated by firms with significant market power, resulting in a relatively small share of surplus accruing to consumers. The decline in welfare contribution we observe in recent years raises important questions about potential barriers to innovation and the changing nature of competition in product markets.

References

Aghion, P., A. Bergeaud, T. Boppart, P. J. Klenow, and H. Li (2019): "Missing growth from creative destruction," *American Economic Review*, 109, 2795–2822.

Amiti, M., M. Gomez, S. H. Kong, and D. E. Weinstein (2024): "Using Stock Returns to Assess the Aggregate Effect of the US-China Trade War,"Technical report, Federal Reserve Bank of New York.

Austin, D. H. (1993): "An event-study approach to measuring innovative output: The case of biotechnology," *American Economic Review*, 83, 253–258.

Balasubramanian, N., and J. Sivadasan (2011): "What Happens When Firms Patent? New Evidence from U.S. Economic Census Data," 93, 126—146.

Bloom, N., M. Schankerman, and J. Van Reenen (2013): "Identifying technology spillovers and product market rivalry," *Econometrica*, 81, 1347–1393.

Bresnahan, T. F., and R. J. Gordon (1996): "The Economics of New Goods,"Technical report, National Bureau of Economic Research.

Broda, C., and D. E. Weinstein (2006): "Globalization and the Gains from Variety," *The Quarterly journal of economics*, 121, 541–585.

——— (2010): "Product creation and destruction: Evidence and price implications," *American Economic Review*, 100, 691–723.

Chaney, P. K., T. M. Devinney, and R. S. Winer (1991): "The Impact of New Product Introductions on the Market Value of Firms," *The Journal of Business*, 64, 573–610.

Chen, S., K. W. Ho, and K. H. Ik (2005): "The Wealth Effect of New Product Introductions on Industry Rivals," *The Journal of Business*, 78, 969–996.

Cohen, L., K. Diether, and C. Malloy (2013): "Misvaluing innovation," *The Review of Financial Studies*, 26, 635–666.

Eddy, A. R., and G. B. Saunders (1980): "New product announcements and stock prices," *Decision sciences*, 11, 90–97.

Ederer, F., and B. Pellegrino (2021): "A tale of two networks: Common ownership and product market rivalry," *SSRN Electronic Journal*.

Eeckhout, J., and L. Veldkamp (2023): "Data and markups: A macro-finance perspective."

Feenstra, R. C. (1994): "New product varieties and the measurement of international prices," *The American Economic Review*, 157–177.

Garcia-Macia, D., C.-T. Hsieh, and P. J. Klenow (2019): "How destructive is innovation?" *Econometrica*, 87, 1507–1541.

Gort, M., and S. Klepper (1982): "Time paths in the diffusion of product innovations," *The economic journal*, 92, 630–653.

Greenstein, S. M. (1996): "From Superminis to Supercomputers: Estimating Surplus in the Computing Market," in *The Economics of New Goods*: University of Chicago Press, 329–372.

Hausman, J. A. (1996): "Valuation of New Goods under Perfect and Imperfect Competition," in *The Economics of New Goods*: University of Chicago Press, 207–248.

Hoberg, G., and G. Phillips (2016): "Text-based network industries and endogenous product differentiation," *Journal of Political Economy*, 124, 1423–1465.

Hopenhayn, H. A., and K. Okumura (2025): "Dynamic Oligopoly and Innovation: A Quantitative Analysis of Technology Spillovers and Product Market Competition," *Working Paper*.

Hsu, P.-H., D. Li, Q. Li, S. H. Teoh, and K. Tseng (2021a): "Valuation of New Trademarks," *Management Science*, (forthcoming).

Hsu, P.-H., K. Li, X. Liu, and H. Wu (2021b): "Consolidating Product Lines via Mergers and Acquisitions: Evidence from the USPTO Trademark Data," *Journal of Financial and Quantitative Analysis*, (forthcoming).

Huber, K. (2023): "Estimating general equilibrium spillovers of large-scale shocks," *The Review of Financial Studies*, 36, 1548–1584.

Kogan, L., D. Papanikolaou, A. Seru, and N. Stoffman (2017): "Technological innovation, resource allocation, and growth," *The Quarterly Journal of Economics*, 132, 665–712.

Nevo, A. (2003): "New products, quality changes, and welfare measures computed from estimated demand systems," *Review of Economics and statistics*, 85, 266–275.

Pakes, A. (1985): "On patents, R&D, and the stock market rate of return," *Journal of Political Economy*, 93, 390–409.

Pellegrino, B. (forth.): "Product Differentiation and Oligopoly: a Network Approach," *American Economic Review*.

Petrin, A. (2002): "Quantifying the benefits of new products: The case of the minivan," *Journal of political Economy*, 110, 705–729.

Schumpeter, J. A. (1943): *Capitalism, Socialism and Democracy*, London: George Allen and Unwin.

Sood, A., and G. J. Tellis (2009): "Do Innovations Really Pay Off? Total Stock Market Returns to Innovation," *Marketing Science*, 28, 442–456.

Wittink, D., A. Ryans, and N. Burrus (1982): "New products and security prices," *Work-ing paper. Ithaca, NY: Cornell University*.

Online Appendix: New Products

Abhiroop Mukherjee HKUST
Bruno Pellegrino Columbia Business School
Alminas Žaldokas National U. of Singapore
Yiman Ren UMichigan Ross
Tomas Thornquist Abu Dhabi Investment Auth.

A Product Survival Probabilities

To estimate the welfare contribution of new products, we need to account for their finite lifespan in the market. Products can be discontinued or become obsolete over time, meaning the profit streams they generate are not perpetual. We estimate the rate at which products exit the market using data from RavenPack.

Using RavenPack's coverage of product mentions in news articles, we can track the survival of products over time. A product is considered "surviving" in a given year if it continues to receive mentions in RavenPack's news coverage. Figure 9 plots the empirical survival probabilities (light blue dots) against a theoretical survival curve (solid blue line) derived from assuming a constant hazard rate λ for product discontinuation.

The empirical pattern of product survival is well approximated by a Poisson process with a constant hazard rate λ= 1.27% per year. As shown in Figure 8, the theoretical survival probability at time t given by $e^{-\lambda t}$ closely tracks the empirical survival curve over a 20-year horizon.

Figure 8: RavenPack: Product Survival Probabilities

