New Products

Abhiroop Mukherjee
HKUST

Bruno Pellegrino
Columbia, NBER

Alminas Žaldokas
National U. of Singapore

Yiman Ren
DRW Trading

Tomas Thornquist
Abu Dhabi Investment Auth.

Abstract

Measuring the welfare impact of new product introductions is a long-standing challenge for economists. In this study, we make progress on this problem by leveraging the informational efficiency of equity markets and a scalable consumer demand model. We construct a novel database of new product announcements covering 20 years (2002-2021) and use stock market reactions to estimate the profits that these new products generate for the inventor firms. We then use the network oligopoly model of Pellegrino (2025) to measure the change in competitors' profits and consumer surplus induced by the new products, thus obtaining the dollar welfare generated. We find that: 1) new products introduced by U.S. publicly-traded firms generate substantial welfare gains, between 1.2 and 2% of US GDP per year; 2) a minority of the announcements account for most of the gains; 3) producer surplus accounts for roughly 40% of these gains. This latter figure is significantly higher than for existing products: we show that this is due to the fact that new product creation is concentrated among firms that have a high degree of market power.

JEL Classification: G14, L1, O3, O4
Keywords: Event Study, Growth, Innovation, Oligopoly, Products

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@nus.edu.sg - 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, Katya Neretina, 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, Society of Economic Dynamics Meeting, European Economic Association Meeting, ESADE Spring Workshop, International Industrial Organization Conference, SGF Conference, Corporate Finance Day, HKUST, 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. Contributions: Abhiroop Mukherjee, Bruno Pellegrino and Alminas Zaldokas are responsible for ideation, planning, methodology, modeling, writing and most of the coding implementation. Former UMichigan graduate student Yiman Ren and former HKUST Assistant Professor Tomas Thornquist provided coding support in the early phases of this project.

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 fully 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 crucial limitation of the IO approach is that it is only feasible to implement in very few industries where the required 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 U.S. publicly-traded firms. We accomplish this with a three-step approach.

The first step consists of constructing a dataset of new product introduction events which spans a 20-year period (2002-2021). Our starting data source is S&P Capital IQ's Key Developments database, which provides a database of time-stamped, product-related announcement events. We use artificial intelligence to single out all those events that can be reliably identified as product introductions. Specifically, we use a state-of-the-art, large language classification model that we previously trained on a human-classified sub-sample.

The second step consists of estimating the yearly profits that these new products generate for the inventor firm. To accomplish this, we rely on the informational efficiency of equity markets and perform a stock market event study: for each new product announcement, we compute the reaction of the inventor firm's stock price to the announcement. The underlying idea behind this approach is that from an asset pricing standpoint, the stock price impact of a media article in a short time-window surrounding its publication reflects the market's perception of the marginal impact of the new product on the firm's future discounted profits. Our approach is similar to Mukherjee, Singh and Žaldokas (2017), and also to Kogan, Papanikolaou, Seru and Stoffman (2017), who value patents using stock market reactions to patent grants. This approach allows us to overcome the challenge that, for most product or product upgrades, we don't observe their contribution to the inventor firm's revenues and costs.

The final step in our methodology is to compute the product market spillovers of new product introductions, and their impact on consumer surplus. In order to accomplish that, we use the network oligopoly model of Pellegrino (2025), which is uniquely suited to this purpose. The network oligopoly model embeds a scalable hedonic demand system (GHL) that can be taken to the data for all U.S. publicly-traded firms. The model allows us to estimate the impact of the product introduction on the focal innovating firm's competitors, and on consumers.

The economic intuition that underlies this third step is that the profits that a new product is expected to generate for the inventor firm reflect two things: 1) the consumers' willingness-to-pay for that good or service, net of the marginal cost required by the firm to produce it; 2) the intensity of product market competition faced by the inventor firm. By isolating the second channel, the network oligopoly model allows us to use our measure of expected profits to pin down the consumer's willingness to pay for the new products.[1]

Using this approach, we find that new products introductions by U.S. public corporations (alone) generate substantial welfare gains, averaging about 300 US$ billion in a typical year – or between 1.2% and 2% of GDP annually.

One obvious limitation of our approach is that our welfare figures only cover product announcements made by publicly traded firms, and not those made by private companies. While public corporations are likely heavily over-represented in the sub-population of firms that routinely introduce or update products, our aggregate measure is still a downward-biased estimate of the economy-wide contribution of new products to welfare growth. In addition, even if we only consider public corporations, these figures could well be a conservative estimate: our welfare metric excludes by construction updates to future profits that are incorporated into stock prices at earlier stages of the innovation process, such as when patents are granted (see Kogan et al., 2017) or when clinical trial results are announced (Rothenstein et al., 2011).

Looking at the cross-section, we find that product innovation generates significant spillovers, both to producers of substitute goods and (to a lesser extent) positive effect for producers of complement goods. A key advantage of our methodology lies in our ability to estimate these spillovers at the level of each firm pair (for a typical year, our dataset contains 15 million firm pairs): this was only previously possible in industry studies.

Our methodology also allows us to investigate how these welfare gains are shared between producers and consumers: we find that about 60% of these gain accrue to the consumer, and the remaining 40% to the innovating firms. For the latter, this is larger than what was previously estimated by Pellegrino (2025) for existing varieties. Further analysis reveals the economic mechanism that underlies this apparent discrepancy: namely, new products are introduced disproportionately more frequently by firms with substantial market power. Because these firms face a less elastic residual demand function, they are able to extract more of the surplus from new and improved products in the form of increased profits.

Our analysis also reveals two key patterns in how innovation benefits are distributed. First, the economic impact of new products is highly concentrated, with the top 10% announcements consistently generating about 90% of both profits and welfare gains. While it is important to clarify that an announcement is not the same as a product (e.g., Apple has over 100 announcements related to the iPhone alone in our dataset) this high level of concentration suggests an key role for "breakthrough" product innovations. Interestingly, we find that this concentration has gradually declined from about 95% in the early 2000s to around 86% by 2021

Second, 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.

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

First, we build on a pre-existing literature in IO that seeks to understand the welfare impact of new products. 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) and Nevo (2003) quantified the

[1] While in principle we could also attempt to measure spillovers to related firms empirically (e.g., see Huber, 2023) —for example, using the competitors' stock price reactions on the day of the product announcement —this approach is likely to suffer from significant measurement error due to the inherent volatility of stock prices. In particular, the further a firm is from the focal innovating firm, the more likely its price response is to be influenced by unrelated factors. For the announcing firm, most of the market-adjusted price change on the announcement day is plausibly attributable to the news, allowing us to use it as a proxy for the value created. In contrast, for less directly related firms, observed price movements are more likely to reflect noise or other confounding influences, making reliable inference more difficult.

welfare benefits of the introduction of minivans and (again) breakfast cereals 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. In a subsequent study, Wittink et al. (1982) focused on computer and office machines businesses. 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 (broadly consistent with our own estimate). 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 utilize stock market reactions to value individual new products, similar to Mukherjee et al. (2017). However, here we take the next step and leverage the structural oligopoly model of Pellegrino (2025) to quantify their impact on aggregate welfare, which is our unique contribution.

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 two important respects. The first key difference is that 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. The second key difference with respect to KPSS lies in the fact that our method allows us to estimate more than just the profits generated by the new products: we but also the consumer surplus and the product market spillovers generated which makes our study unique in this literature.

Pakes (1985) also provides an early contribution examining the relation between patents and the stock market rate of return. Aside from KPSS and Pakes (1985), 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 estimate 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 total factor productivity (TFP) growth into various components (own innovation, creative destruction, and the introduction of new varieties) using a model of growth with creative destruction. Finally, Boppart, Klenow, Laski and Li (2025) build further on this agenda, showing how to use corporate valuations to infer innovation step size. 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.

A separate literature exists in applied macroeconomics that investigates product entry, exit and life cycles using micro

data (in particular, the Nielsen Retail Scanner Data): examples include Argente, Baslandze, Hanley and Moreira (2020), Babus, Moreira and Marzani (2023) and Argente, Lee and Moreira (2024). While we connect to this literature in that we are interested in product entry, we differ from it by virtue of the source of micro data used (we obtain our data from product announcements in Capital IQ, and therefore span a larger set of industries) as well as the underlying research question: in this study we focus on the novel challenge of quantifying the welfare contribution of new products, as opposed to studying their life cycle dynamics.

This study contributes to an emerging literature that leverages large network oligopoly models (Pellegrino, 2025; Ederer and Pellegrino, 2021; Eeckhout and Veldkamp, 2023; Hopenhayn and Okumura, 2025); 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 on the econometric identification of spillovers (see Bloom, Schankerman and Van Reenen, 2013 and Huber, 2023 among others). In addition, we note the previous work of Hoberg and Phillips (2010), who show that mergers between firms with complementary product-market positions lead to post-merger scope expansion. In their setting, new product creation is proxied by growth in the length and content of product descriptions in 10-K filings, whereas in our setting we observe new products directly through their announcements in Capital IQ.

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 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 Building the Dataset of Product Announcements

Source Data. Our source of new product announcements data is S&P's Capital IQ Key Developments database. This database tracks significant corporate events, including new product introductions, across publicly traded companies, and offers several advantages for our purposes.

First, the Capital IQ Key Developments feed is curated by S&P Global Market Intelligence specifically to capture firm-level announcements that are likely to be material for investors—for example, news that could plausibly move a firm's stock price. Each entry is timestamped to reflect the moment when the information becomes public, based on S&P's aggregation of corporate press releases, regulatory filings, and major newswire services. This focus on material, market-relevant disclosures makes the resulting timestamps particularly suitable for event-study analysis.

A second key advantage of Capital IQ Key Developments is its extensive coverage: it covers the universe of publicly-traded firms over more than two decades (since the late 1990s).

A third advantage of this data source is that it provides structured metadata —including standardized event type codes— which facilitates the identification of new product introduction events.

Due to these features, Capital IQ's Key Developments has been utilized extensively in the finance literature for studying a variety of events that move stock prices (e.g., Lee et al., 2018). The presence of significant abnormal returns around these announcements in prior work supports the interpretation that Key Developments entries correspond to genuine market-moving news. Together, these features make the dataset a reliable foundation for constructing our panel of new-product announcements.

From this starting sample, we construct a novel firm-day panel of verified new product announcements. The goal is to isolate product news that reflects meaningful innovation from unrelated or confounding corporate events, and to do so at scale with high precision. This section outlines the step-by-step methodology used to build this measure, including raw data collection, filtering, classification, and validation.

Refinement Step A: Initial Sample and Confounder Filtering. The starting sample for our analysis is a comprehensive set of product-related corporate announcements from Capital IQ, covering the period 2002 through 2021. We begin with 1,128,419 raw observations that we obtain by focusing on event type code number 41 ("Product-Related Announcements"), which captures various forms of product-related news including product launches, unveilings, and major enhancements. However, this tag is also inclusive of a wide range of content, including marketing campaigns, product launch anniversaries, or pricing changes – which are not about new product launches.

To ensure that we isolate the effect of product announcements from other events that also move stock prices, we drop any headline that occurs within a three-day window (t – 1 to t + 1) around major corporate announcements. Specifically, we exclude windows centered on: 1) Earnings announcements, sourced from Compustat; 2) Mergers and acquisitions (M&A), from SDC Platinum, and 3) Dividend declarations and stock splits, from Capital IQ's corporate actions dataset.

This step reduces the likelihood that observed stock returns are driven by overlapping events rather than product news.

Later in the construction process, to preserve aggregate value estimates, we impute the value of announcements dropped due to confounding windows. Specifically, we assume the lost announcements would, on average, resemble those in the clean sample in terms of their average return effect.

Refinement Step B: Data Cleaning and De-duplication. From the raw Capital IQ set, we first remove entries that lack a valid firm identifier (gvkey), which are necessary for linking announcements to firm-level financial data. We also eliminate duplicate entries, defined by the tuple (headline, announcement date, gvkey). That is, if the same firm is assigned the same headline on the same day more than once, we retain only one copy.

Note that headlines involving multi-firm collaborations (e.g., joint ventures or co-branded products) are retained as separate rows for each participating firm. This reflects the fact that such announcements may carry firm-specific market implications and are treated accordingly.

After cleaning and de-duplication, we are left with a refined dataset of 645,529 firm-headline-date observations.

Refinement Step C: Manual Labeling of Training Set. Next, we utilize human-supervised artificial intelligence (AI) to assess how many of the headlines in the dataset thus obtained actually represent genuine new product launches. We begin by drawing a random sample of 10,002 headlines, stratified by year to ensure temporal coverage across the 2002–2021 period.

Each headline in this sample is manually reviewed by a human research assistant and assigned a binary label: "1" if it reflects a true new product announcement (e.g., a firm launches a new device, software, drug, platform, service); "0" if it does not (e.g., refers to a product introduced in the past, a feature expansion, marketing language, or a product-related milestone). This manually-classified sub-sample forms the training dataset for the artificial intelligence that is utilized in the subsequent step.

The manual classification exercise reveals that only 47% of product-tagged headlines are true new product announcements, with the remaining 53% being mis-classified or ambiguous. This highlights the need for a more sophisticated filtering method than simple keyword matching or dictionary-based classification.

Refinement Step D: AI Classification. Because the classification problem utilizes short text and is likely to be context-sensitive, traditional text analysis tools (e.g., keyword filters or bag-of-words models) are likely perform poorly at this task. Slight variations in phrasing—such as "launches" versus "updates," or "integration" versus "release"—can flip the label. We therefore turn to modern natural language understanding (NLU) tools, specifically the BERT (Bidirectional Encoder Representations from Transformers) family of models. Specifically, we use DeBERTaV3-Large, a state-of-the-art language model developed by Microsoft. BERT models differ fundamentally from generative. In contrast to generative large language models (LLMs) such as ChatGPT (which are designed to produce open-ended text), BERT models are optimized to extract meaning from short pieces of language—such as business headlines—by predicting relationships between words and recognizing nuanced context. This makes this class of model particularly well suited to our task of identifying which corporate announcements correspond to genuine new-product launches.

To make the model task-specific, we fine-tune it using our manually-labeled sample of 10,002 observations. This involves three steps: A) splitting the sample into 80% training, 10% validation, and 10% test sets; B) training the model for 10 epochs; 3) Selecting the best-performing checkpoint, which achieves 96.4% accuracy on the out-of-sample test set. The fine-tuned model is then used as our classifier for the remainder of the product announcement dataset.

Refinement Step E: Classification of the Full Universe. We apply the trained model to the full cleaned universe of 645,529 Capital IQ product-tagged announcements. Each headline is assigned a binary label indicating whether it is a verified new product announcement. The model classifies 312,020 of these firm-headline observations as true new product launches, yielding our final dataset.

To ensure consistency, if a headline in the inference set exactly matches one from the manually labeled sample but receives a different model label, we overwrite the model prediction with the manual label. This guarantees that all manually reviewed headlines are correctly classified.

The model's positive prediction rate is 48%, closely matching the 47% positive rate in the labeled sample, suggesting generalization across time and firm types.

Advantages of the Approach. This human-supervised machine learning classification strategy offers three key advantages. The first such advantage is precision at scale: the event-study analysis we conduct below relies on precise identification of genuine product news. Misclassifying generic or retrospective announcements as new products would attenuate measured stock reactions. Fine-tuning a large Transformer model on domain-specific labels enables high precision and recall across hundreds of thousands of observations, far outperforming dictionary-based or rule-based methods.

The second advantage is domain robustness and temporal stability: headlines span a 27-year period, during which corporate language and communication styles have evolved. BERT-family models, pre-trained on large corpora and fine-tuned on a stratified sample, generalize well across time and firm types. This reduces time-varying misclassification that could otherwise bias trend estimates or risk models.

The third advantage is clean integration: our AI classifier produces a simple binary label that directly reflects our conceptual definition of "new product launches." This label can be seamlessly integrated into our empirical pipeline, including event studies and survival analysis, to estimate the economic value of innovation announcements.

We provide a more complete discussion on our data cleaning procedures in Appendix A.

2.3 Estimating the Profit Impact of New Products

In the next step of our quantification exercise, we leverage the efficient market hypothesis (Fama, 1970) to obtain a forward-looking measure of the yearly profits generated by each new product announcement. The key idea underlying this step is that –to the extent that a product announcement leads investors to revise their expectations of the inventor company's future profits– the magnitude of the stock market reaction to the announcement should reflect the incremental profits generated by the product.

For our quantification exercise, we define the stock market reaction as the 1-day abnormal return on the inventor firm's stock. Expected returns are calculated based on the Fama-French 3-factor model.

From Abnormal Returns to Expected Yearly Profits. Our event study yields an estimate of the percent change in a firm's market capitalization as it announces a new or improved product. However, what we need for our welfare measurement exercise is a measure of yearly profits that are generated by the corresponding product.

The two can be naturally connected using a discounted utility model: through the lens of a discounted utility model, in fact, a firm's market capitalization must equal the expected present value of all its future cashflows. It follows that the change in a company's market valuation that is directly attributable to the announcement of a new product ought to equal the expected present value of future profits generated by newly-announced product.

Because different term structures of future profits can produce the same stock market reaction, in order to pin down a unique measure of future profits we need to make an assumption: we assume that every new product z, introduced by a generic firm i, generates a constant stream of profits v_{iz} until it is discontinued. The yearly hazard rate at which the product is discontinued is also constant and equal to h.

Under this assumption, and by applying the perpetuity present value formula, the yearly contribution of new product z to the profits of its inventor firm i can be estimated as:

$$v_{iz} \;=\; \left[r_{it(z)} + h\right]\; \delta_{iz}\, \mathcal{M}_{it(z)-1} \tag{2.1}$$

where δ_{iz} is the percent increase in the market capitalization of company i caused by the product introduction z; $r_{it(z)}$ is the discount rate applicable to firm i's equity cashflows at time $t\,(z)$ – that is, the day on which new product z is announced; $\mathcal{M}_{it(z)-1}$ is the previous day's market capitalization.

In our quantification exercise, we calibrate h to a value of 12% per year, consistent with a mean survival of 8.3 years and a median survival of 5.8 years. We base this calibration on the product life cycle analysis of Argente, Lee and Moreira (2024), who estimate a Kaplan Meier survival function for products included in the Nielsen Consumer Retail Scan dataset: their analysis is consistent with the hazard rate of discontinuation progressively decreasing over the life cycle of products. This means that a higher hazard rate ($\sim$ 28% per year) is required to match survival rates over very short horizons, while at longer horizons such as 10 years a lower hazard rate ($\sim$ 8% per year) is required. We select a value of h closer to the end of this range, both to be conservative (a higher h implies higher expected profits conditional on survival given the effect on market cap) and because our expected profitability measures (obtained market cap) are by construction more likely to capture the long-term profitability of products.

Bayesian Filtering of the Abnormal Returns. Now, it may be tempting to use the measured abnormal return on the date of the announcement - which we call $\mathcal{A}_{it(z)}$ - to proxy for δ_{iz}. However, this would lead us to infer that the profit impact is negative for nearly half of the product announcements in our dataset. This is an artifact of the fact that 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 or idiosyncratic volatility.

FIGURE 1: PRODUCT ANNOUNCEMENTS: FILTERED ABNORMAL RETURNS

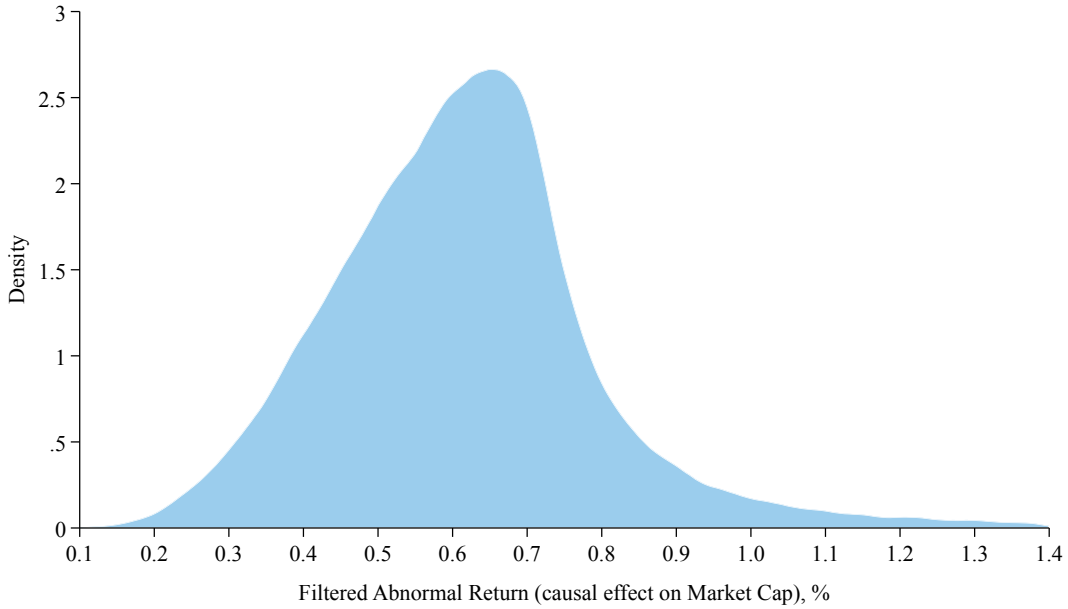


FIGURE NOTES: This figure shows the kernel density estimate of the distribution of filtered abnormal returns. Each observation is a firm-day when a new product is announced. The sample includes all new product announcements from Capital IQ that were identified by our methodology and that did not coincide with other major corporate events (earnings announcements, M&A, or dividend/split announcements) in a three-day window around the announcement.

In the next step we follow KPSS's approach and use a Bayesian technique to filter out the noise. We follow KPSS in assuming that the abnormal return on the announcement date is the sum of δ_{iz} and an idiosyncratic error term $\varepsilon_{it(z)}$

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

where the signal follows a normal distribution that is left-truncated at zero, with underlying mean μ_δ and standard deviation parameter τ_δ; the noise follows a normal distribution with mean zero and time-varying volatility τ_{it}^2:

$$\delta_{iz} \;\sim\; \text{Trunc}_{[0,\infty)}\, \mathcal{N}\left(\mu_\delta, \tau_\delta^2\right) \qquad \varepsilon_{it} \;\sim\; \mathcal{N}\left(0, \tau_{it}^2\right) \tag{2.3}$$

This specification is similar to that used by KPSS, except that it is more general: KPSS assume a half-normal distribution for the signal (effectively, they impose $\mu_\delta = 0$). Using the properties of the truncated normal distribution, we can derive the expectation of the abnormal return along with the expected squared return:

$$\mathbb{E}\left[\mathcal{A}_{it(z)}\right] \;=\; \mu_\delta + \tau_\delta \cdot \frac{\varphi\left(\mu_\delta/\tau_\delta\right)}{\Phi\left(\mu_\delta/\tau_\delta\right)} \tag{2.4}$$

$$\mathbb{E}\left[\mathcal{A}_{it(z)}^2\right] \;=\; \tau_\delta^2 \left\{1 - \frac{\mu_\delta}{\tau_\delta} \cdot \frac{\varphi\left(\mu_\delta/\tau_\delta\right)}{\Phi\left(\mu_\delta/\tau_\delta\right)} - \left[\frac{\varphi\left(\mu_\delta/\tau_\delta\right)}{\Phi\left(\mu_\delta/\tau_\delta\right)}\right]^2\right\} + \left[\mu_\delta + \tau_\delta \cdot \frac{\varphi\left(\mu_\delta/\tau_\delta\right)}{\Phi\left(\mu_\delta/\tau_\delta\right)}\right]^2 + \tau_{it}^2 \tag{2.5}$$

where $\varphi(\cdot)$ and $\Phi(\cdot)$ denote, respectively, the Gaussian probability density function and cumulative distribution function. Because all terms of equations (2.4) and (2.5) are observed, except for μ_δ and τ_δ, we can straightforwardly estimate these two parameters using a simple GMM procedure. In this case, because there are two equations and two unknowns, the system is exactly identified.

Again, using the properties of the truncated normal distribution, we obtain the following posterior expectation for δ_{iz},

FIGURE 2: PRODUCT ANNOUNCEMENTS: SIGNAL-TO-NOISE RATIO

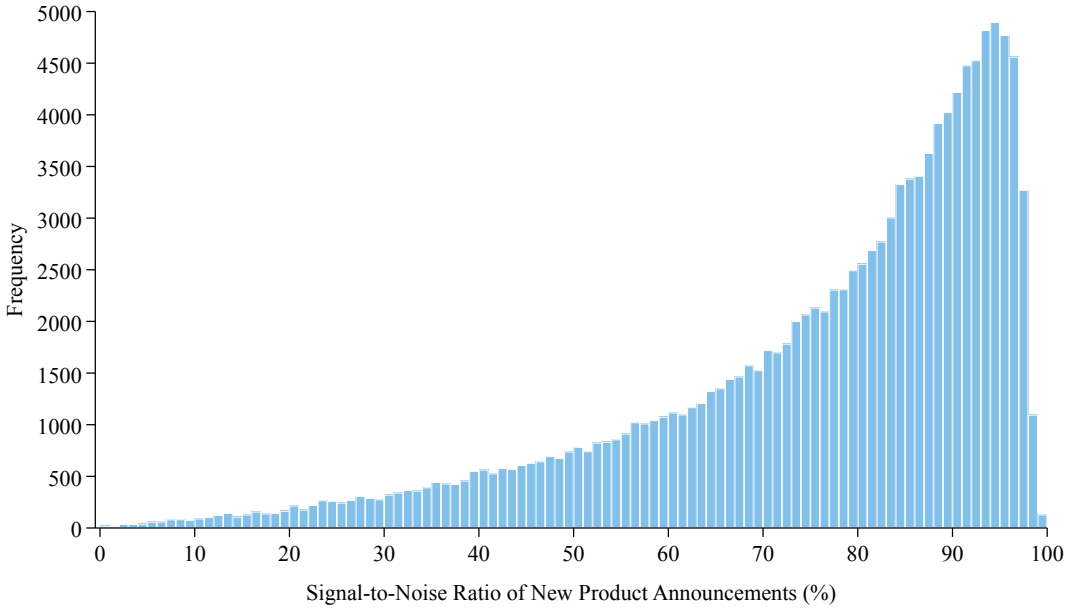


FIGURE NOTES: The figure above is a histogram that represents the distribution of the signal-to-noise ratios, defined in equation (2.8). Each observation is a firm-day when a new product is announced.

conditional on the observed abnormal return $\mathcal{A}_{it(z)}$:

$$\mathbb{E}\left[\delta_{iz}|\,\mathcal{A}_{it(z)}\right] \;=\; \tilde{\mu}_{it(z)} + \tilde{\tau}_{it(z)} \cdot \frac{\varphi\left(\tilde{\mu}_{it(z)}/\tilde{\tau}_{it(z)}\right)}{\Phi\left(\tilde{\mu}_{it(z)}/\tilde{\tau}_{it(z)}\right)} \tag{2.6}$$

where $\tilde{\tau}_t$ and $\tilde{\mu}_t$ are defined as follows:

$$\tilde{\tau}_{it(z)} \overset{\text{def}}{=} \sqrt{\frac{1}{\tau_\delta^{-2} + \tau_{it}^{-2}}}; \qquad \tilde{\mu}_{it(z)} \overset{\text{def}}{=} \frac{\tilde{\tau}_{it(z)}^2}{\tau_\delta^2}\mu_\delta \;+\; \frac{\tilde{\tau}_{it(z)}^2}{\tau_{it}^2}\mathcal{A}_{it(z)} \tag{2.7}$$

We use the expectation in equation (2.6) as our estimate of δ_{iz}, and plug it inside equation (2.1) to obtain our estimate of v_{iz}.

Figure 1 plots the distribution of the filtered abnormal returns $\mathbb{E}\left(\delta_{iz}|\,\mathcal{A}_{it(z)}\right)$. The distribution is bell shaped and centered around 63 basis points. It features a pronounced positive skew and right fat tail, consistent with occasional large positive market reactions to particularly impactful announcements. Most filtered returns falling between 30 and 90 basis points.

One reason why our filtered returns are so heterogeneous is that the majority of the announcements in our dataset generate clean abnormal returns that are contaminated by little noise. We can measure the lack of noise in the data using the Signal-to-Noise Ratio (SNR), which in our case is time-varying and given by:

$$\text{SNR}_{iz} \;=\; \left\{1 + \frac{\tau_{it}^2}{\tau_\delta^2}\left[1 - \frac{\mu_\delta}{\tau_\delta}\cdot\frac{\varphi\left(\mu_\delta/\tau_\delta\right)}{\Phi\left(\mu_\delta/\tau_\delta\right)} - \left(\frac{\varphi\left(\mu_\delta/\tau_\delta\right)}{\Phi\left(\mu_\delta/\tau_\delta\right)}\right)^2\right]^{-1}\right\}^{-1} \tag{2.8}$$

Figure 2 presents the distribution of SNR across all our events (firm-days). We find that the vast majority of the events have a very high SNR: nearly 90% have an SNR above 50%. The average is 76% and the median is 82%. By

Figure 3: Product Announcements - Expected Profits

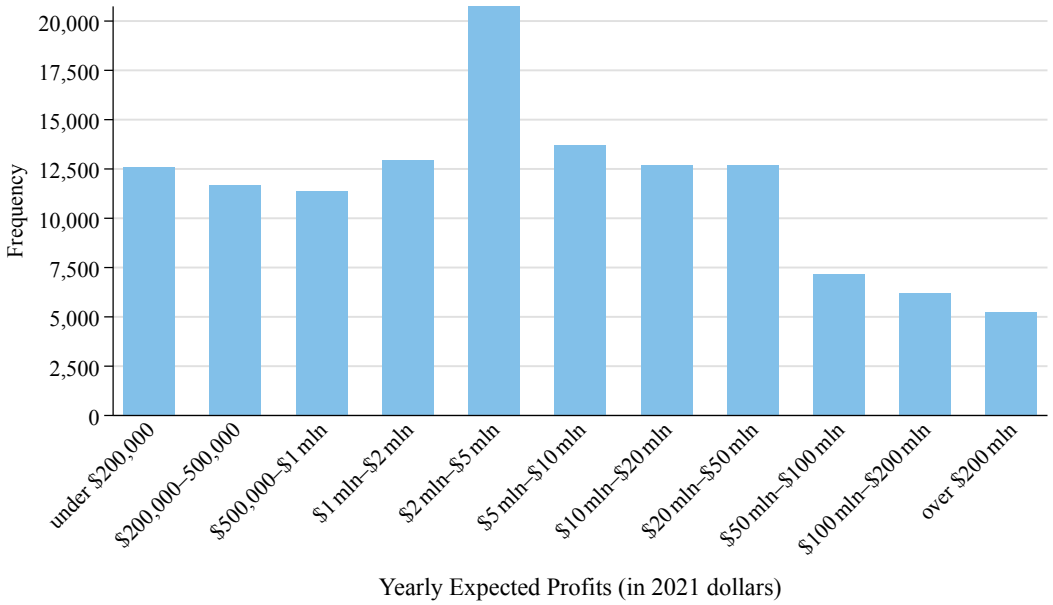


Figure Notes: This figure shows the distribution of the estimated yearly profits (in 2021 U.S. dollars) attributed to new product announcements. Each observation is a firm-day when a new product is announced. The sample includes all new product announcements from Capital IQ that were identified by our 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.

comparison, the (constant) SNR in the dataset of KPSS is 1.4%.

Figure 3 presents the distribution of estimated profit dollar values implied by these announcement returns. The distribution is centered between $2–$5 million, with a long right tail extending to values above $200 million, illustrating the positive skew in firm-level innovation outcomes.

Positive Filtered Returns and Free-Disposal Assumption. As in KPSS, this filtering procedure yields expected profits (v_{iz}) that are non-negative by construction. We acknowledge that there is disagreement in the literature over the merits of this specific filtering approach: some researchers[2] believe that this approach artificially "discards" unsuccessful product introductions – i.e. the introduction of products that prove ex-post to be commercial failures, and on which firms never manage to break even (e.g., Google Glass, Microsoft Zune,...).

We thus want to take the opportunity to discuss why, while we acknowledge this critique, we feel we can safely reject it and defend KPSS's filtering method on a logical basis. Negative new product profit contributions are unrealistic in our setting (as in KPSS) because they are inconsistent (at least over the medium-long term) with a free disposal assumption. In other words, even when a product "bombs", the inventor firm typically learns quickly that consumers have a low willingness to pay for the product (it has a clear incentive to do so), and swiftly discontinues the unsuccessful product: this makes large "true" negative reactions unlikely: a short-lived sequence of negative unit profits combined a low volume is unlikely to have a large negative impact on the inventor firm's equity value.

Methodological Differences with respect to Kogan et al. (2017). Overall, our valuation method is similar to that utilized by KPSS to value patents, except for a few key differences. The first, as we saw earlier, is to use a more general specification for the distribution of the signals, which allows us to untie the mean of δ_{iz} from its variance. A key reason for this choice is that by tying the dispersion of the signal to the mean, the half-normal assumption would lead us to

[2]We thank in particular our early conference discussants Grzegorz Pawlina and Katya Neretina who brought this debate to our attention.

Table 1: Summary Statistics of New Products

	Abnormal Return	**Filtered AR**	**Expected Profits**
Statistic	$\mathcal{A}_{it(z)}$ - (%)	$\mathbb{E}\left[\delta_{iz} \middle\vert \mathcal{A}_{it(z)}\right]$ - (%)	v_{jz} - ($mln)
Mean	0.378	0.636	34.078
Standard Deviation	6.364	0.305	115.336
Percentiles			
- $p01$	-7.639	0.249	0.029
- $p05$	-3.602	0.342	0.095
- $p10$	-2.362	0.398	0.202
- $p25$	-0.988	0.500	0.791
- $p50$	0.018	0.613	3.827
- $p75$	1.120	0.710	19.468
- $p90$	2.873	0.830	87.667
- $p95$	4.736	0.963	176.207
- $p99$	12.464	1.573	429.536

Table Notes: This table presents summary statistics for new product announcements by U.S. publicly traded firms (2002–2021). "Abnormal Return" ($A_{it(z)}$) is the daily abnormal stock return (in %), estimated using the Fama–French 3-factor model. "Filtered AR" (δ_z) is the abnormal return after correcting for measurement error (following Kogan et al., 2017) or alternatively the causal effect of the product introduction on the inventor firm's market capitalization. "Expected Profits" (v_{jz}) is the implied annual profit stream (in millions of dollars), derived from equation (2.1). Mean, standard deviation, and selected percentiles ($p1, p5, \ldots, p99$) are reported. Monetary values are in 2021 dollars and returns are expressed in percentage points.

under-estimate our signal-to-noise ratio by multiple orders of magnitude, leading to excessive shrinkage and a resulting loss of valuable information that comes from the abnormal returns themselves. By relaxing the assumption $\mu_\delta = 0$ we can obtain an accurate estimate of our SNR. This in turn allows us to utilize efficiently the information contained in our abnormal returns.

A second methodological difference with respect to KPSS's study lies in the fact that, thanks to the nature of the underlying data, we are able to use a shorter (1-day) event window than KPSS (3-day).

A third important difference lies in our assumptions about the signal-to-noise ratio (SNR). KPSS assume a constant SNR for abnormal returns across firms and over time. As noted by Federle, Harhoff and Kreyer (2025) this assumption, while convenient (and thus justifiable in the context of KPSS's exercise) has one key disadvantage: it implies that the signal component is higher in times of higher idiosyncratic volatility. But because the half-normality assumption ties the mean of the signal distribution to the variance, this leads to the implication that the *mean* of the signal is higher in times of higher idiosyncratic volatility. As a result, this set of assumptions (which we also used in a preliminary version of this paper) produces innovation series that are very highly correlated with aggregate stock market volatility.

For this reason, in equations (2.3)-(2.7) we departed from KPSS and did *not* assume a constant SNR. Instead, we held the signal variance constant across firms and time and let the idiosyncratic error variance vary over time (equation 2.3). This prevents the expected value of new product announcements from being mechanically related to market volatility (and thus keeps our aggregate innovation series from spuriously comoving with aggregate stock market volatility).

2.4 Summary Statistics

Table 1 provides comprehensive summary statistics on new product announcements by U.S. publicly traded firms from 2002-2021. The table presents three key metrics: raw abnormal returns $\left(\mathcal{A}_{it(z)}\right)$, filtered abnormal returns after accounting for measurement error (δ_z), and the implied annual profit streams (v_{jz}) derived from these announcements.

The raw abnormal returns average 0.378%, but exhibit substantial variability with a standard deviation of 6.364%. The distribution is slightly positively skewed, with a median of just 0.018%. The 10th percentile shows negative returns of -2.362%, while the 90th percentile shows positive returns of 2.873%, illustrating the wide range of market reactions to new product announcements.

After applying the measurement error correction methodology described in equation (2.6), the filtered abnormal returns (v_{jz}) show a much tighter distribution. The mean filtered return is 0.636% with a considerably smaller standard deviation of 0.305%. The range from the 1st to 99th percentile spans from 0.249% to 1.573%, representing the market's estimated true valuation effect of these product introductions.

The expected profits (v_{jz}), calculated using equation (2.1), average \$34.078 million per new product, though the distribution is highly skewed with a standard deviation of \$115.336 million. The median annual profit contribution is just \$3.827 million, while the 90th percentile reaches \$87.667 million and the 99th percentile is close to \$429.536 million. This substantial right skew in the profit distribution highlights that a small percentage of breakthrough products generate disproportionately large economic impacts.

These statistics underscore the highly concentrated nature of innovation benefits, where the top new products account for the vast majority of the profits generated.

3 Network Oligopoly: Spillovers and Welfare

3.1 Consumer 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 network oligopoly model of Pellegrino (2025), which we review briefly below. It is a highly tractable can be solved analytically and taken to the data for thousands of firms and millions of firm pairs.

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 and a unit-length vector $\mathbf{a}_i \in \mathbb{R}^n$ of common characteristics, formally:

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

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

These characteristic (column) vectors can be stacked next to each other inside a matrix that we call $\mathbf{A}_t$. Let q_{it} be the quantity of good i consumed by the representative agent at time t. The total amount of characteristic k consumed by the agent at time t is then given by

$$x_{kt} = \sum_{i=1}^{n} a_{ki} q_i \tag{3.3}$$

We assume that there is a representative consumer, whose utility $\mathcal{U}$ is a quadratic function of the vector of common characteristics $\mathbf{x}_t$ and the vector of idiosyncratic characteristics, and is linear (decreasing) in the labor supply, L_t:

$$\mathcal{U}(\mathbf{x}_t, \mathbf{q}_t, L_t) \quad \stackrel{\text{def}}{=} \quad \sum_{i=1}^{n} b_{it} q_{it} - \frac{\alpha}{2} \cdot \sum_{k=1}^{m} x_{kt}^2 - \frac{1-\alpha}{2} \sum_{i=1}^{n} q_{it}^2 - L_t \tag{3.4}$$

where b_{it} is the intercept demand shifter for product i and α is the weight placed by the consumer on the common characteristics. Because each unit of good i provides exactly one unit of its corresponding idiosyncratic characteristic, we simply use q_{it} (the quantity consumed of good i) denote the units of idiosyncratic characteristic i consumed by representative consumer.

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. Pellegrino (2025) shows that the following linear demand system solves the consumer problem:

$$\mathbf{p}_t \;=\; \mathbf{b}_t - (\mathbf{I} + \mathbf{\Sigma}_t)\, \mathbf{q}_t \tag{3.5}$$

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

$$\mathbf{\Sigma}_t \quad \stackrel{\text{def}}{=} \quad \frac{\partial \mathbf{p}_t}{\partial \mathbf{q}_t} \quad = \quad \alpha \left(\mathbf{A}_t' \mathbf{A}_t - \mathbf{I} \right) \tag{3.6}$$

as explained in Pellegrino (2025) is a matrix that captures substitution among products. 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. Geometrically, it measures the cosine of the angle formed by vectors $\mathbf{a}_{it}$ and $\mathbf{a}_{jt}$ in the space of characteristics. The higher the cosine similarity $\mathbf{a}_{it}' \mathbf{a}_{jt}$, the more similar the characteristic content of the products sold by firms i and j . A high cosine similarity between goods i and j implies that the products of firms i and j are highly substitutable.

An important advantage of GHL demand is that it allows for rich substitution patterns, and does not rely on an arbitrary delineation of industries or markets. The model of oligopoly Pellegrino (2025) that is built on this demand system is described in the literature as a "network" model because firms are akin to nodes in a network where the links represent cosine similarity between goods. The centrality of a firm reflects its exposure to intense product market competition, and determines the elasticity of demand faced by the firm. Firms whose products are highly central have many competitors who produce products with similar characteristics: these firms have less market power. Vice-versa, firms who are peripheral sell products with unique characteristics, and therefore are more akin to monopolists.

3.2 Equilibrium Profits and Welfare

Each firm i faces a marginal cost of production equal to c_{it}, and maximizes a profit function $\pi_{it}(\mathbf{q})$, defined as follows:

$$\pi_{it}(\mathbf{q}_t) \quad \stackrel{\text{def}}{=} \quad [p_{it}(\mathbf{q}_t) - c_{it}]\, q_{it} - f_{it} \tag{3.7}$$

where f_{it} is the fixed cost of operating firm i at time t, which we assume to be sunk (we do not model entry). Similarly, we can define the economic profits $\pi_{it}^E(\mathbf{q})$, which are simply profits gross of fixed costs:

$$\pi_{it}^E(\mathbf{q}_t) \quad \stackrel{\text{def}}{=} \quad [p_{it}(\mathbf{q}_t) - c_{it}]\, q_{it} \tag{3.8}$$

In the network oligopoly model of Pellegrino (2025), firms compete in quantities (Cournot competition) and choose their own output q_{it} to maximize profits. The vector of economic profits $\boldsymbol{\pi}_t^E$ that obtains at the Cournot-Nash equilibrium is given by:

$$\boldsymbol{\pi}_t^E \quad = \quad \mathbf{\Omega}_t (\mathbf{b}_t - \mathbf{c}_t) \tag{3.9}$$

where the matrix $\mathbf{\Omega}_t$ is defined as follows:

$$\mathbf{\Omega}_t \quad \stackrel{\text{def}}{=} \quad (2\mathbf{I} + \mathbf{\Sigma}_t)^{-1} \tag{3.10}$$

and $\mathbf{b}_t$ and $\mathbf{c}_t$ are, respectively, column vectors. Each row of $\mathbf{b}_t$ is the demand function intercept of a product and the corresponding entry of $\mathbf{c}_t$ represents the marginal cost incurred to produce one unit. There is a closed-form expression for the equilibrium Consumer Surplus, which we denote by S_t:

$$S_t \;=\; \frac{1}{2} \cdot (\mathbf{b}_t - \mathbf{c}_t)' \, \mathbf{\Lambda}_t \, (\mathbf{b}_t - \mathbf{c}_t) \tag{3.11}$$

where the matrix $\mathbf{\Lambda}_t$ is defined as follows:

$$\mathbf{\Lambda}_t \;\stackrel{\text{def}}{=}\; \mathbf{\Omega}_t \;(\mathbf{I} + \mathbf{\Sigma}_t) \;\mathbf{\Omega}_t \tag{3.12}$$

3.3 Measuring Product Market Spillovers

We now consider how product introductions can change equilibrium profits and consumer surplus. While we already have analytical expressions for $\boldsymbol{\pi}_t$ and S_t, the tractability of the framework presented in the previous subsection will also allow us to derive closed-form expressions for the firm-to-firm spillover that generated by the introduction of a new product.

By total differentiation of equation (3.9), the change in profit from $t-1$ to t is equal (to a first-order approximation) to:

$$\Delta\pi_{it} \;\approx\; 2\sqrt{\pi^e_{jt}} \left[\sum_{j=1}^{n} \omega_{ijt} \, \Delta\left(b_{jt} - c_{jt}\right) \;+\; \sum_{j=1}^{n} \left(b_{jt} - c_{jt}\right) \, \Delta\omega_{ijt} \right] - \Delta f_{it} \tag{3.13}$$

where Δ indicates the 1-period difference. The first term in parentheses reflects parallel shifts in willingness to pay and marginal costs (vertical innovations). The second reflects the effect of changes in the characteristics content of products (horizontal innovations).

Assumption: Vertical Innovations. In order to compute explicit formulas for the product market spillovers generated by product introductions, we make an important assumption: namely, we assume that the introduction of a new product contained in our dataset only affects equilibrium profits and consumer surplus through changes in $b_{jt} - c_{jt}$ and changes in fixed costs (f_{it}). This assumption effectively requires that the cosine similarity $\mathbf{A}'\mathbf{A}$ to be stable over the time interval considered for the measurement (one year) or at least that its innovations are independent of $\mathbf{b}$. 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.

It is reasonable to question how restrictive this assumption is: after all, firms often introduce products with different characteristics mix from existing products. We make the case that this assumption is reasonable in our specific setting due to the nature of our dataset. Crucially, the database that we created does not contain products, but rather *product announcements.* Some of the most valuable products in our database (e.g., Microsoft Windows, Apple iPhone, etc...) are associated with dozens, possibly hundreds of announcements: they do not reach their full profitability potential the moment they are introduced, but rather they are improved announcement after announcement. Over small horizons, it is reasonable to assume that the characteristics content of these goods stays approximately the same.

To model the impact of new product introductions, we assume that changes in $(b_{it} - c_{it})$ from period to period can be broken down into two components:

$$\Delta\left(b_{jt} - c_{jt}\right) \;=\; \sum_{z \in \mathcal{Z}_{jt}} \beta_{jz} + \xi_{jt} \tag{3.14}$$

The first is a summation over $\mathcal{Z}_{jt}$ be the set of new products introduced by firm i, between $t-1$ and t. Each term of the summation, β_{jz} corresponds to one of the products introduced by firm i between times $t-1$ and t, and it captures the effect of introducing z on $\left(b_{jt} - c_{jt}\right)$. The second component, ξ_{jt}, is a random shock which captures all other economic forces affecting $\left(b_{jt} - c_{jt}\right)$, such as costs shocks.

Similarly, we break down changes in fixed costs f_{jt} into a component that captures the effect of new products, and a residual component ζ_{jt}:

$$\Delta f_{jt} \;=\; \sum_{z \in \mathcal{Z}_{jt}} \theta_{jz} + \zeta_{jt} \tag{3.15}$$

where θ_{zt} is the impact of product z, introduced by firm j at time t, on the fixed costs of inventor company j. Combining equations (3.13)-(3.15), we obtain an expression for v_{jz} – the effect of product z's introduction on firm j's profits:

$$v_{jz} \quad \stackrel{\text{def}}{=} \quad \frac{\partial \pi_{jt}}{\partial b_{jt}} \cdot \beta_{zt} \quad = \quad 2\sqrt{\pi^E_{jt}} \cdot \omega_{jjt} \cdot \beta_{zt} - \theta_{jz} \tag{3.16}$$

Fixed Costs Incidence. In section 2, we saw how we could obtain an estimate of v_{jz} from observing the reaction of the price of company j's stock to the announcement of new product z. The next step consists of using the network oligopoly model to recover β_{zt} from v_{jz}. One obstacle that we face in doing so is that we do not observe directly θ_{jz}, the fixed costs impact of product z: therefore, we shall estimate it from aggregate data. Specifically, we assume that, for all product announcements z in our dataset, fixed costs account for a fixed fraction $\mathcal{F}$ of the total economic profits generated by z – formally:

$$\frac{\theta_{jz}}{v_{jz} + \theta_{jz}} \;=\; \mathcal{F} \quad \forall \quad (j, z) \tag{3.17}$$

To estimate $\mathcal{F}$ we map, at the firm/year level, Selling, General and Administrative (SGA) to fixed costs (f_{it}) and Costs of Goods Sold (COGS) to total variable costs ($c_{it} q_{it}$) as in Pellegrino (2025). We then estimate $\mathcal{F}$ as the aggregate incidence of SGA on operating margins – formally:

$$\mathcal{F} \;=\; \frac{\sum_{t=2002}^{2021} \; \sum_{i=1}^{n} \; (f_{it})}{\sum_{t=2002}^{2021} \; \sum_{i=1}^{n} \; (p_{it} q_{it} - c_{it} q_{it})} \tag{3.18}$$

Armed with the fixed costs share $\mathcal{F}$, we can now proceed to derive an explicit expression for the product market spillovers and the consumer surplus effect of z.

Spillovers to Peers. We define v_{iz}, the effect of the introduction of z on a generic company i (not necessarily equal to j, who introduced the product) as follows:

$$v_{iz} \quad \stackrel{\text{def}}{=} \quad \frac{\partial \pi_{it}}{\partial b_{jt}} \cdot \beta_{zt} \quad = \quad 2\,(1 - \mathcal{F})\,\sqrt{\pi^E_{it}} \cdot \omega_{ijt} \cdot \beta_{jz} \tag{3.19}$$

Then, combining equations (3.16) and (3.19) we can finally write the spillover to firm j in terms of the private rent v_{jz} and other observables ($\boldsymbol{\pi}$ and $\boldsymbol{\Omega}$):

$$v_{iz} \quad = \quad \sqrt{\frac{\pi^E_{it}}{\pi^E_{jt}}} \cdot \frac{\omega_{ijt}}{\omega_{jjt}} \cdot \frac{v_{jz}}{1 - \mathcal{F}} \tag{3.20}$$

Consumer Surplus. Finally, we consider the effect on consumer surplus S_t, which we denote with the capital letter V:

$$V_z \quad \stackrel{\text{def}}{=} \quad \frac{\sum_{i=1}^{n} (b_{it} - c_{it})\,\lambda_{ijt}}{2\sqrt{\pi^E_{jt}} \cdot \omega_{jjt}} \cdot \frac{v_{jz}}{1 - \mathcal{F}} \tag{3.21}$$

Note that the effect on the full consumer surplus includes both the direct effect from focal new product as well as effects on consumer surplus coming from complementary and substitute spillovers on other firms, e.g., competing or complementary products.

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 introduction of new products.

FIGURE 4: VALIDATION USING TRADEMARKS AND RAVENPACK PRODUCTS

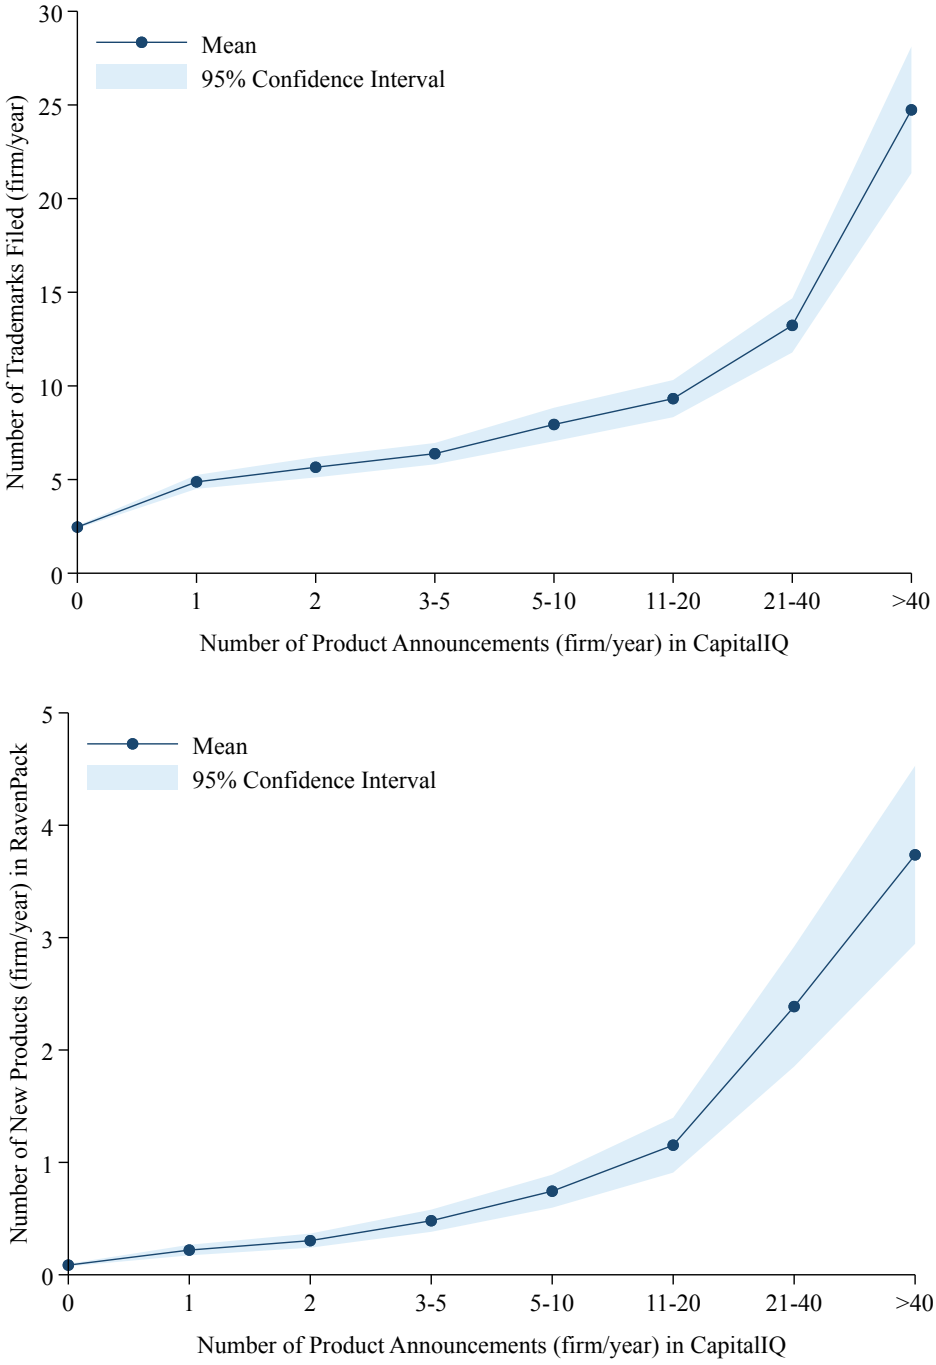


FIGURE NOTES: The figures above shows the relationship between product announcements in Capital IQ and trademarks (upper figure) and product announcements in Capital IQ and product introductions identified from RavenPack dats (lower figure) over 2002-2019. The horizontal axis shows bins for the number 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.

3.4 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 (2025) 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): this dataset, constructed from similarities between textual descriptions of products contained in the mandatory annual 10-K filings of the public corporations, provides a time-varying empirical estimate of $\mathbf{A}_t'\mathbf{A}_t$. 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 are filed with US Patent and Trademark Office (USPTO) when a firm launches a new product line or service (Hsu, Li, Liu and Wu, 2021b; Hsu, Li, Li, Teoh and Tseng, 2021a).[3] The upper figure in Figure 4 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 4 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.

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.

[3] We thank the authors of Hsu, Li, Liu and Wu (2021b); Hsu, Li, Li, Teoh and Tseng (2021a) for sharing this data.

TABLE 2: NEW PRODUCT INTRODUCTIONS VS. PATENTING

		Receives a Patent		Total
		No	Yes	
Announces New Products	No	3,730 (39.5%)	748 (7.9%)	4,478 (47.4%)
	Yes	2,153 (22.8%)	2,824 (29.9%)	4,977 (52.6%)
Total		5,883 (62.2%)	3,572 (37.8%)	9,455 (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 9,455 unique firms. "Announces New Products" is a dummy variable equal to one if a firm has at least one new product announcement in the Capital IQ 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.

4.2 New Product Introductions vs. Patents

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 2 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 41.2%), there is substantial divergence between these two measures of innovation.

Perhaps the most notable finding is that 22.8% 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 technologies.

At the same time, 7.9% 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 (39.5%) consists of firms that neither patent nor introduce new products during this period.

Notably, 29.9% of firms both receive patents and introduce new products, indicating that these two innovation metrics often complement each other. However, the fact that over 40% of the innovating firms in our sample (22.8% out of 52.6%) 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 an-

FIGURE 5: AGGREGATE WELFARE CONTRIBUTION OF NEW PRODUCTS (2001-2021)

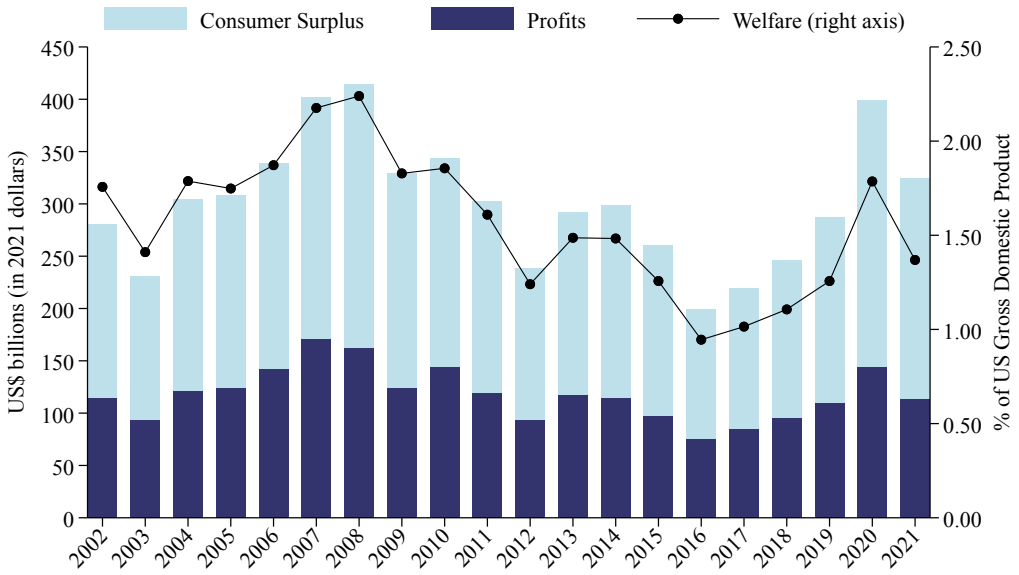


FIGURE NOTES: This figure shows the aggregate welfare contribution of new product introductions by publicly-traded U.S. firms from 2002 to 2021. The stacked bars represent the composition of welfare gains, decomposed into producer surplus (profits, shown in dark blue) and consumer surplus (shown in light blue), measured in billions of 2021 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.

nouncement measure thus provides important complementary information about innovation activity, particularly in sectors where patenting may be less relevant or strategic secrecy more valuable.

5 Results

5.1 Welfare Contribution of New Products

Figure 5 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 about 1% to 2.3% of GDP over the sample period. The welfare contribution peaked in 2008, when new products generated approximately $150 billion in profits and $250 billion in consumer surplus, totaling about 2.3% of GDP. Throughout our sample period, the distribution of welfare gains remains relatively stable, with consumer surplus consistently accounting for roughly 60-65% of the total benefits, while producer surplus comprises the remaining 35-40%. It may be worth reiterating here that these are gains from a subset of product announcements – our method applies only to publicly traded firms and does not include products introduced by private firms; and even for these firms, they reflect a subset of the total gains from innovation – we do not account for expectations about profit changes that the market might have already impounded into stock prices arise at earlier stages of innovation, e.g., when patents are granted or when clinical trial success is announced.

FIGURE 6: AGGREGATE WELFARE DECOMPOSITION (2001-2022)

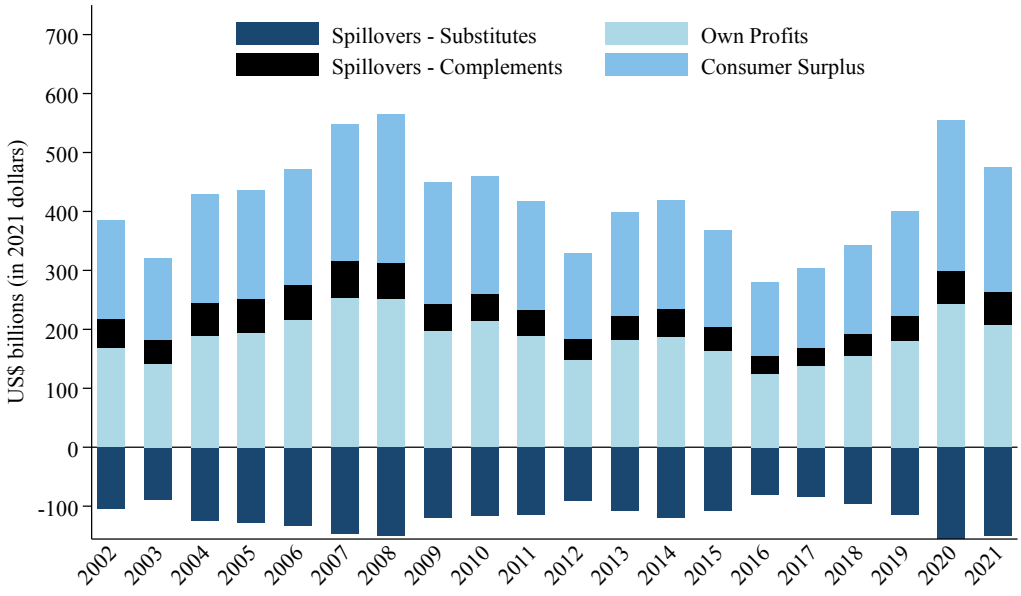


FIGURE NOTES: this figure decomposes the aggregate welfare contribution of new product introductions by U.S. publicly-traded firms into four components over the period 2002-2021. The stacked bars show: (1) consumer surplus (mid blue), representing the direct benefit to consumers; (2) complementary profits (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 2021 U.S. dollars.

By the end of our sample period in 2021, new products generated approximately $100 billion in profits and $200 billion in consumer surplus, amounting to about 1.8% 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 6 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 a substantial share of welfare gains, averaging around $120-250 billion annually during our sample period. However, these gains are partially offset by negative spillovers to competitors through substitution effects, which typically amount to $80-150 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 $20-50 billion annually throughout most of the sample period.

Consumer surplus remains a significant and stable component of total welfare, typically ranging between $130-250 billion annually. This relatively stable share of consumer surplus suggests that firms have maintained consistent pricing

FIGURE 7: HETEROGENEITY OF SPILLOVERS TO COMPETITORS AND CONSUMER (2021)

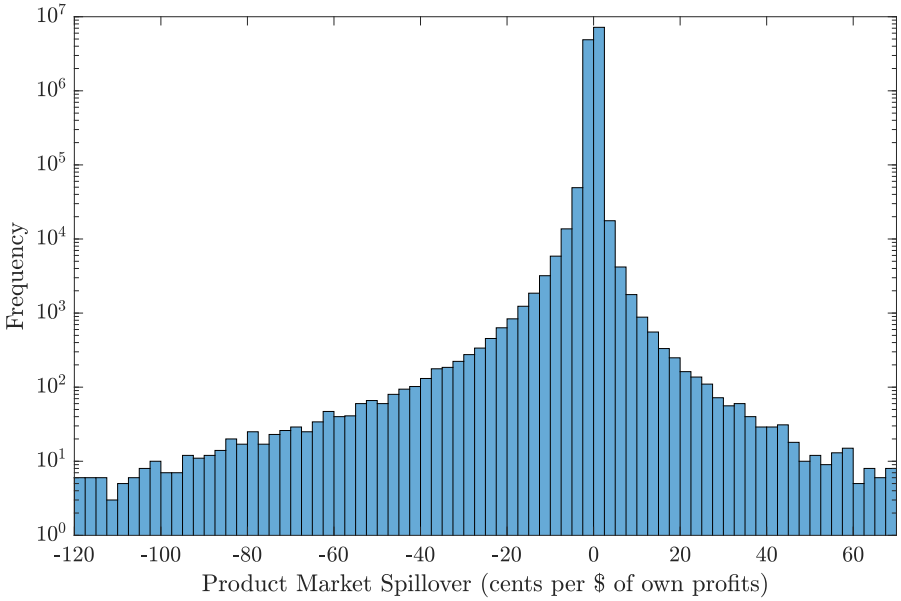


FIGURE NOTES: this figure shows the distribution of product market spillovers generated by new product introductions in 2021, measured in cents per dollar of the innovating firm's own profits, while the vertical axis shows the frequency on a logarithmic scale. Spillovers are estimated using our GHL demand 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.

power over their new products throughout our sample period, despite changes in market structure and competitive dynamics.

Figure 7 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 8 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.

FIGURE 8: VALUE OF NEW PRODUCTS: CONCENTRATION

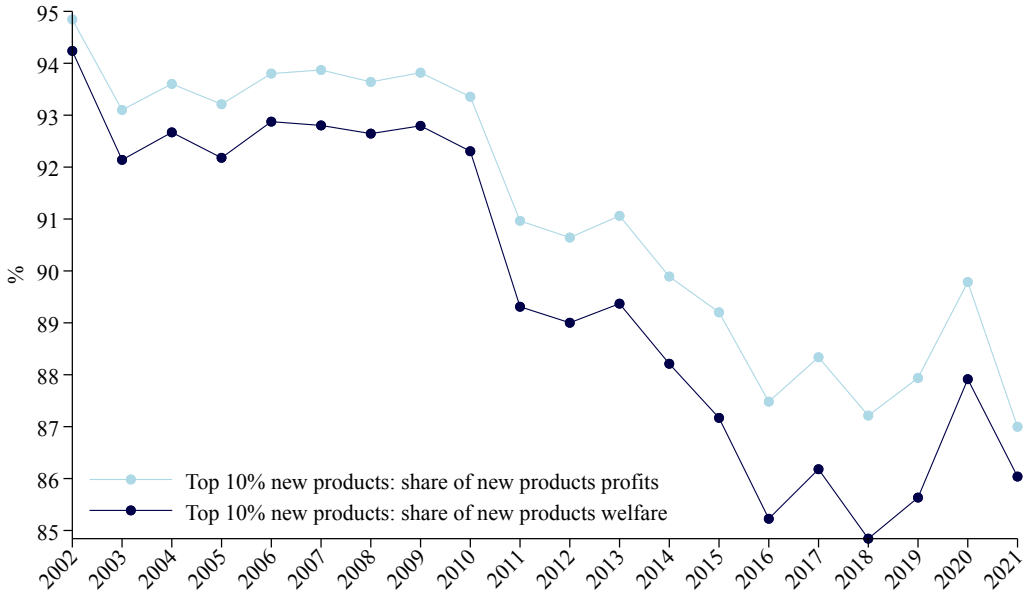


FIGURE NOTES: This figure shows the share of total profits and welfare generated by the top 10% of new product introductions by U.S. publicly-traded firms over the period 2002-2021. The light blue line represents the share of total new product profits accounted for by the top decile of products (ranked 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.

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 90% in 2015 to 88% by 2021. The concentration of welfare effects shows an even more pronounced decline, dropping to approximately 87% 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.

The declining concentration over time could reflect several factors, including 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 smaller 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'

Figure 9: Distribution of Product Market Centrality

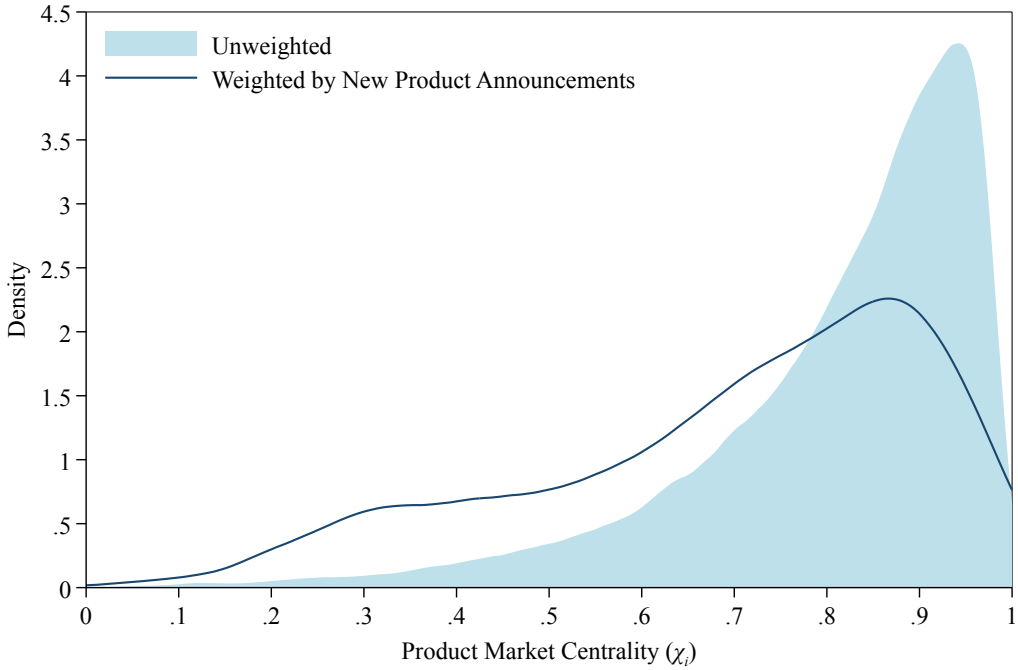


Figure Notes: This figure shows two distributions of product market centrality (χ_i) for U.S. publicly-traded firms in 2021. Product market centrality measures a firm's competitiveness based on its position in the network of product market rivalries, ranging from 0 (monopolist-like behavior) to 1 (perfect 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.

market power, we employ the measure of product market centrality (χ_i) developed by Pellegrino (2025) – 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 9 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) 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 $\chi_i = 0.2 - 0.8$. 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 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 1-2.3% of GDP annually during our sample period. These gains are split between producer surplus (roughly 40%) and consumer surplus (roughly 60%), suggesting that firms are able to capture a significant portion of the value they create through innovation. This relatively smaller 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, 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 86% of both profits and total welfare benefits. While this concentration has gradually declined from about 94% 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 after 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. 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.

Argente, D., S. Baslandze, D. Hanley, and S. Moreira (2020): "Patents to products: Product innovation and firm dynamics."

Argente, D., M. Lee, and S. Moreira (2024): "The life cycle of products: Evidence and implications," *Journal of Political Economy*, 132, 337–390.

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

Babus, A., S. Moreira, and M. Marzani (2023): "The rise of specialized financial products," *Available at SSRN 5575993*.

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

Boppart, T., P. J. Klenow, R. Laski, and H. Li (2025): "Idea Rents and Firm Growth," *Unpublished Manuscript*.

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."

Fama, E. F. (1970): "Efficient capital markets: A review of theory and empirical work," *The journal of Finance*, 25, 383–417.

Federle, J., D. Harhoff, and A.-C. Kreyer (2025): "Value Estimates for U.S. Patent Grants," February, `https://www.ifw-kiel.de`, Preliminary version.

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.

He, P., J. Gao, and W. Chen (2021): "Debertav3: Improving deberta using electra-style pre-training with gradient-disentangled embedding sharing," *arXiv preprint arXiv:2111.09543*.

Hoberg, G., and G. Phillips (2010): "Product market synergies and competition in mergers and acquisitions: A text-based analysis," *The Review of Financial Studies*, 23, 3773–3811.

——— (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.

Lee, J., A. Naranjo, and G. Velioglu (2018): "When do CDS spreads lead? Rating events, private entities, and firm-specific information flows," *Journal of Financial Economics*, 130, 556–578.

Mukherjee, A., M. Singh, and A. Žaldokas (2017): "Do corporate taxes hinder innovation?" *Journal of Financial Economics*, 124, 195 – 221.

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. (2025): "Product differentiation and oligopoly: A network approach," *American Economic Review*, 115, 1170–1225.

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

Rothenstein, J. M., G. Tomlinson, I. F. Tannock, and A. S. Detsky (2011): "Company stock prices before and after public announcements related to oncology drugs," *Journal of the National Cancer Institute*, 103, 1507–1512.

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 DRW Trading
Tomas Thornquist Abu Dhabi Investment Auth.

A Additional Details on the New Products Detection Algorithm

We begin with a dataset of product-related announcements from the WRDS Capital IQ - Key Developments database and focus on event type 41 (Product-Related Announcements), which includes product launches, unveilings, and major enhancements. The dataset includes variables such as headline, announcement date, and company identifiers. In cases where headlines describe multiple companies collaborating on a product launch, multiple observations are present with the same headline and date but different company identifiers corresponding to the companies involved.

To isolate the stock-market effect of product news from other corporate events, we exclude announcements that coincide with major confounders: earnings announcements (Compustat), M&A (SDC Platinum), and dividend declarations or stock splits (WRDS Capital IQ - Key Developments: Dividends and Splits). To preserve aggregate value estimates, we later impute and add back the value of new products that would have come from dropped windows by assuming they resemble the average clean-sample announcement.

Our initial Capital IQ product-related set spans 1997–2023 with 1,128,419 observations. We remove entries missing a company identifier (gvkey) and drop duplicates defined by the tuple (headline, announcement date, gvkey). Headlines mentioning multi-firm collaborations remain as separate firm-level rows by construction.

Next, we assess the accuracy of Capital IQ's new product announcement classification for our specific purpose, which is to identify firm announcements about new product launches. To do so, we first draw a random sample of 10,002 observations stratified by year, and manually add a binary label indicating whether the headline actually pertains to a new product announcement or is mislabeled (for example, referring to a clinical trial update of a pharmaceutical drug or a drilling update by a natural resource exploration company). This manual classification finds that approximately 47% of headlines are true new product announcements, and 53% are not, highlighting the need for filtering.

This is a short-text, context-sensitive classification problem: small changes in wording ("launches," "unveils," "updates," "beta," "pilot," "feature expansion," "integration," etc.) often flip the label, reducing the effectiveness of traditional word-list-based textual analysis. Modern machine learning models, on the other hand, are purpose-built for such natural language understanding (NLU) tasks because they learn rich, context-dependent representations of text and transfer effectively to downstream classifiers with limited labeled data. Here we choose the BERT (Bidirectional Encoder Representations from Transformers) class of models. This is an encoder-only Transformer that uses self-attention to encode every word in the context of all other words in the sequence. It is first pre-trained on large unlabeled corpora with self-supervised objectives (originally masked language modeling), which teaches the model general linguistic and semantic patterns.

We use DeBERTaV3-Large, a modern BERT-family encoder that improves upon classic BERT/DeBERTa in two key ways that matter for our purpose: (1) Pre-training objective: Instead of masked language modeling, DeBERTaV3 adopts a detection method which is more sample-efficient and yields stronger NLU features for the same compute budget. (2) Training stability and embedding quality: The model introduces gradient-disentangled embedding sharing

to avoid the "tug-of-war" between generator and discriminator embeddings during training, improving both efficiency and downstream performance. These two improvements deliver strong results across standard NLU benchmarks (e.g., GLUE), with DeBERTaV3-Large setting new state-of-the-art performance among models of similar structure He, Gao and Chen (2021). We leverage these gains because headline classification hinges on fine-grained semantic distinctions, domain phrasing, and long-range dependencies even in short text.

For a specific task like headline classification, DeBERTaV3-Large has to be "fine-tuned". For this purpose, the manually labeled random sample of 10,002 observations above is added on top of the encoder's pooled representation (i.e, the pre-trained model), allowing the "machine" to learn about our specific context. This fine tuning yields state-of-the-art accuracy on sentence- and document-level classification, including short, noisy business headlines.

Specifically, for fine-tuning, we split the 10,002 manually labeled observations into three subsets: 80% for training, 10% for validation, and 10% for out-of-sample testing. Our random sample contains no duplicate headlines, thereby eliminating concerns about data leakage between the three subsets. Training and validation are performed for 10 epochs, and the checkpoint with the highest accuracy is selected as the best fine-tuned model. When evaluated on the out-of-sample test set, it achieves 96.3% accuracy. We then use this model for headline classification in the inference stage.

The final step is to extend inference on the full universe. Because Capital IQ is updated from time to time, we re-download the most up-to-date product-related universe at the time of writing (1997–2023; 1,131,623 observations). After cleaning for missing gvkeys and removing duplicates in terms of headline, gvkey, announcement date and time, 645,402 observations remain. We apply the fine-tuned DeBERTaV3-Large model to classify each headline. The positive rate in inference is ≈49%, closely matching the 47% positive rate in the labeled random sample mentioned above. When an inference headline exactly matches one from the labeled sample but receives a different model-predicted label, we overwrite the model prediction with the manual label. We retain observations classified as new product announcements, yielding 316,067 firm-headline observations.