New Products

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

Abstract

Quantifying the contribution of new product introductions to economic growth is a major, long-standing challenge in economics. This paper develops a measure of the welfare contribution of new products introduced by US public corporations, constructed from a three-step approach. One, we use convolutional neural networks (CNNs) to identify media articles about new product launches. Two, we use the stock market reaction to the publication of these articles to estimate the present value of the future profits generated by the new products. Three, drawing from modern empirical IO techniques, we use a large and scalable demand system (GHL) to back out, via closed-formed expressions, the change in competitors' profits and consumer welfare induced by these new product introductions, and finally obtain (after aggregating across firms) a measure of the corresponding change in aggregate welfare, expressed in dollars.

1. Introduction

Technological innovation is commonly divided into two types: process innovation (improvements in production technology that allow to use inputs more efficiently) and product innovation (the introduction of new products). Economists developed, several decades ago, the conceptual toolbox that allows us today to measure process innovation at the aggregate level using Total Factor Productivity (TFP), which can in turn be inferred from national accounts under well-established assumptions.

The measurement of product innovation has, however, remained much more elusive. It is not captured by real GDP growth (which captures instead the volume growth of existing products). It is also difficult to map new product introductions to changes in welfare, since the contribution that new products bring to welfare depends on their degree of similarity/substitutability with existing products.

Some progress has been made, at the micro level, by empirical industrial organization economists, using the tools of demand modeling. Estimating demand allows IO economists to recover, from observed price, quantity sold, and product characteristics data, the welfare contribution of new products. The main limitation of this approach is that it is only feasible for very few industries where this rich data is available. In other words, it's not scalable. At the aggregate level, we have instead measures of the economic impact of new patented technologies - a related but distinct concept. Quantifying the value of new product introductions at scale remains a largely open research question.

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

2. Methodology

2.1. New Product Introductions

The first step in our methodology is to identify new product introductions across a wide range of industries and firms. To accomplish this, we develop a systematic approach using convolutional neural networks (CNNs) to analyze media articles from Dow Jones' Factiva database.

We begin by extracting all media articles from Factiva that involve firms traded on US stock exchanges and are filed under the New Products/Services category over the period 1989-2019. This initial dataset comprises 660,958 articles involving 16,278 distinct firms. To ensure relevance, we filter these articles to keep only those where the listed firm appears within the title or the first 50 words, confirming that the product refers to that specific firm. This filtering process reduces our dataset to 326,398 articles.

The next step is to classify these articles accurately. We employ a convolutional neural network (CNN) to distinguish between articles that truly are first mentions of new products and those that are not. To train the CNN, we utilize a custom-built visual interface in the form of a mobile app, through which undergraduate students labeled 2,000 articles each. Each article is assigned to two separate students, and we keep only the articles where both students agree on their classifications.

This process results in a final training set of 15,160 labeled articles, of which 3,762 are judged to be truly about new product or service mentions. We then use Google's pre-trained Word2Vec word embeddings to classify the remaining articles in our sample. The final k-fold out-of-sample results give us a precision (ratio of true positives) of 93% and a recall (ratio of positive articles found) of 86%, resulting in an F1 score of 89%. This high accuracy ensures that our measure of new product introductions is reliable and comprehensive.

2.2. Estimating the Profit Impact of New Products

The second major step in our methodology is to estimate the profit impact of the new product introductions we've identified. Our goal here is to quantify the expected profit impact of new products as perceived by the market, providing a forward-looking measure of the value created by product innovations. Our approach is inspired by and builds upon the methodology developed by Kogan, Papanikolaou, Seru and Stoffman (2017) (hereafter KPSS) for valuing patents.

We start by estimating abnormal returns on the mentioned firm's stock surrounding the day of the article's release. Specifically, we calculate cumulative abnormal returns over a two-day window (0,1) after the announcement. Expected returns are calculated based on the market model. This method allows us to capture the market's immediate reaction to the new product announcement, which should reflect the expected present value of future profits generated by the new product.

To isolate the effect of the new product announcement, we exclude days if the firm in the article announced earnings or an M&A transaction on that day (including one day before and one day after for both events). We then filter our results to focus only on positive abnormal returns, ensuring that we capture products that are important enough for the firm to have a meaningful impact on its value.

Using this data, we create two measures of New Products at the firm-year level. The first is a count of new products launched by a firm in a given year, defined as the number of distinct times we uncover news media mentions of new products associated with the firm that lead to positive abnormal returns. The second measure aggregates cumulative abnormal returns (conditional on them being positive) coming from all new product news for a given firm in a given year.

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

capitalization. This gives us an estimate of the market value added for each new product announcement.

It's important to note that while our approach is inspired by KPSS, there are key differences. We focus on new products, which are the final outputs of innovation that directly reach consumers. This allows us to capture innovation that might not be patentable or where firms choose not to patent for strategic reasons.

By leveraging the efficient market hypothesis in this way, we obtain a forward-looking measure of the private value of new products for every firm/year.

2.3. Demand and Welfare

The third step is to estimate the product market spillovers and welfare contribution of new product introductions. In order to accomplish that, we need a demand system that can cover all US publicly-traded firms. The standard and natural tool in this case is the Generalized Hedonic-Linear (GHL) demand system of Pellegrino (forth.), which we review briefly below.

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 $i \in \{1, 2, ..., n\}$. We assume that each unit of product i provides: 1) one unit of the corresponding idiosyncratic characteristic i; 2) a unit-length vector of $\mathbf{a}_i \in \mathbb{R}^n$ of common characteristics, formally:

$$\mathbf{a}_i \;=\; \left[\begin{array}{cccc} a_{1i} & a_{2i} & \dots & a_{ni} \end{array}\right]' \tag{2.1}$$

$$\text{such that} \qquad \sum_{k=1}^{m} a_{ki}^2 = 1 \quad \forall\, i \in \{1, 2, ..., n\} \tag{2.2}$$

With regard to the n idiosyncratic characteristics, we assume that each unit of good i provides exactly one unit of its corresponding idiosyncratic characteristic i. Hence, we can just use q_i to denote the units of idiosyncratic characteristic i consumed by representative consumer.

We assume the representative consumer's utility function $\mathcal{U}$ is linear quadratic in both common characteristics x_j and idiosyncratic characteristics vector q_i, and linear (decreasing) in the labor supply, L:

$$\mathcal{U}(\mathbf{x}, \mathbf{q}, H) \quad \stackrel{\text{def}}{=} \quad b_i q_i - \frac{\alpha}{2} \cdot \sum_{k=1}^{m} x_k^2 - \frac{1-\alpha}{2} \sum_{i=1}^{n} q_i^2 - L \tag{2.3}$$

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

$$\mathbf{p} \;=\; \mathbf{b} - (\mathbf{I} + \mathbf{\Sigma})\, \mathbf{q} \tag{2.4}$$

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

$$\mathbf{\Sigma} \quad \stackrel{\text{def}}{=} \quad \frac{\partial \mathbf{p}}{\partial \mathbf{q}} \quad = \quad \alpha\,(\mathbf{A}'\mathbf{A} - \mathbf{I}) \tag{2.5}$$

as explained in Pellegrino (forth.) the (i, j) entry of matrix $\mathbf{A}'\mathbf{A}$, equal to $\mathbf{a}_i'\mathbf{a}_j$, is called the "cosine similarity" between firm i and firm j, and is a measure of product substitution. The Cournot-Nash equilibrium vector

of profits $\boldsymbol{\pi}$ is given by:

$$\sqrt{\boldsymbol{\pi}} \quad = \quad \boldsymbol{\Omega}\,(\mathbf{b}-\mathbf{c}) \tag{2.6}$$

$$\text{where} \qquad \boldsymbol{\Omega} \quad \stackrel{\text{def}}{=} \quad (2\mathbf{I}+\boldsymbol{\Sigma})^{-1} \tag{2.7}$$

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

$$S \;=\; \frac{1}{2}\cdot(\mathbf{b}-\mathbf{c})'\,\boldsymbol{\Lambda}\,(\mathbf{b}-\mathbf{c}) \tag{2.8}$$

where

$$\boldsymbol{\Lambda} \;\stackrel{\text{def}}{=}\; \boldsymbol{\Omega}\,(\mathbf{I}+\boldsymbol{\Sigma})\,\boldsymbol{\Omega} \tag{2.9}$$

2.4. Data for GHL demand estimation

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

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

2.5. New Products

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

$$\Delta\pi_{it} \;\approx\; 2\sqrt{\pi_{it}}\left(\sum_{j=1}^{n}\omega_{ijt}\,\Delta b_{jt} \;+\; \sum_{j=1}^{n} b_{jt}\,\Delta\omega_{ijt}\right) \tag{2.10}$$

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

We make an important assumption here, which could be restrictive but that we believe we can verify empirically. Namely, we assume that, the second term in parentheses is small:

$$\sum_{j=1}^{n} b_{jt}\,\Delta\omega_{ijt} \;\approx\; 0 \tag{2.11}$$

this assumption requires that the cosine similarity $\mathbf{A}'\mathbf{A}$ is approximately constant over the time interval considered for the measurement (one year) and/or that its increments are broadly uncorrelated with $\mathbf{b}$. What this assumption means, intuitively, is that we can model new product introductions, over short time periods, as vertical innovations – which affect the demand intercept $\mathbf{b}$ but not $\mathbf{A}'\mathbf{A}$. In other words, over small (1-year) time horizons, new products tend to be upgraded versions of existing products.

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

$$\Delta\pi_{it} \;\approx\; 2\sqrt{\pi_{it}}\sum_{j=1}^{n}\omega_{ijt}\,\Delta b_{jt} \tag{2.12}$$

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

Let us index new products introduced by firm i with the generic subscript z. Let also $\mathcal{Z}_{it}$ be the set of new products introduced by firm ii, between $t-1$ and t. We model changes in b_i over time, as the sum of the effect of the introduction of new products (ξ_{zt}) and of a random shock ε_{it} (the underlying process is irrelevant for our measurement):

$$\Delta b_{it} \;=\; \sum_{z \in \mathcal{Z}_{it}} \xi_{zt} + \varepsilon_{it} \tag{2.13}$$

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

$$r_{iz} \quad \stackrel{\text{def}}{=} \quad \frac{\partial \pi_{it}}{\partial b_{it}} \cdot \xi_{it} \quad = \quad 2\sqrt{\pi_{it}} \cdot \omega_{iit} \cdot \xi_{zt} \tag{2.14}$$

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

2.6. Spillovers and Welfare Impact

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

$$r_{jz} \quad \stackrel{\text{def}}{=} \quad \frac{\partial \pi_{jt}}{\partial b_{it}} \cdot \xi_{zt} \quad = \quad 2\sqrt{\pi_{jt}} \cdot \omega_{jit} \cdot \xi_{zt} \tag{2.15}$$

Then, combining equations (2.14) and (2.15) we can finally write the spillover to firm j in terms of the private rent ξ_i and other observables ($\boldsymbol{\pi}$ and $\boldsymbol{\Omega}$):

$$r_{jz} \quad = \quad \sqrt{\frac{\pi_{jt}}{\pi_{it}}} \cdot \frac{\omega_{jit}}{\omega_{iit}} \cdot r_{iz} \tag{2.16}$$

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

$$R_z \quad \stackrel{\text{def}}{=} \quad \frac{\partial S_t}{\partial b_{it}} \cdot \xi_{iz} \quad = \quad \frac{\sum_{j=1}^{n} \lambda_{ijt} \left(b_{jt} - c_{jt}\right)}{2\sqrt{\pi_{it}} \cdot \omega_{iit}} \cdot r_{iz} \tag{2.17}$$

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

2.7. Defending the Key Assumption

To defend the key assumption (2.11), we plan to use the fact that by taking the GHL demand system to the data, we observe the matrix $\boldsymbol{\Sigma}_t$ and can recover the demand intercept b_{it} and marginal cost c_{it} for each firm. As a consequence, we can compute both terms in equation (2.10) and compare their relative size.

References

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

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

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