

Quality Predictability and the Welfare Benefits from New Products: Evidence from the Digitization of Recorded Music

Luis Aguiar

European Commission

Joel Waldfogel

University of Minnesota and National Bureau of Economic Research

We explore the consequence of quality unpredictability for the welfare benefit of new products, using recent developments in recorded music as our context. We quantify the effects of new music on welfare using an explicit structural model of demand and entry with potentially unpredictable product quality. On the basis of plausible forecasting models of expected appeal, a tripling of the choice set according to expected quality adds substantially more consumer surplus as the usual long-tail benefits from a tripling of the choice set according to realized quality.

I. Introduction

The rise of the internet—and digitization more generally—has placed attention on the welfare benefit of cost reductions that raise the number of

We are grateful for comments from Avi Goldfarb and Ajay Agrawal, as well as seminar participants at the Institute for Prospective Technological Studies, the University of Minnesota, the Society for Economic Research on Copyright Issues, the Summer Institute on Competitive Strategy at Berkeley, the NBER Summer Institute joint Industrial Organization/Digitization meeting, the Econometric Society meetings in Minneapolis, the 12th Conference on Media Economics in Naples, the 13th Centre for European Economic Research (ZEW) Conference on Information and Communication Technologies in Mannheim, the 8th ICT Conference in ParisTech, Penn State, the University of Michigan, Yale University, Stanford Graduate School of Business, the University of Zurich, and Universidad Carlos III de Ma-

Electronically published March 12, 2018

[*Journal of Political Economy*, 2018, vol. 126, no. 2]

© 2018 by The University of Chicago. All rights reserved. 0022-3808/2018/12602-0004\$10.00

available products. Researchers and others have viewed the internet as delivering infinite shelf space, allowing consumers access to a long tail of obscure products.¹ Despite the importance of long-tail effects in consumption, the welfare benefit of new products is much larger when we account for the unpredictability of product quality at the time of investment. We term this the “long tail in production.”

The usual long tail idea in consumption is that the internet allows consumers access to the large number of extant products rather than simply the popular products that consumers might access from a local retailer with limited shelf space. While access to additional products is clearly beneficial to consumers, the benefits may be somewhat limited: given the substitutability among differentiated products, the incremental benefit of obscure products—even lots of them—can be small. A long tail in production is different. The appeal of many products to consumers is difficult to know at the time that investments are made. This unpredictability is substantial for cultural products such as books, movies, and music, leading screenwriter William Goldman (1984) to famously remark that “nobody knows anything” about which new movies will be commercially successful. Industry observers report that roughly 10 percent of new movies are commercially successful, and the figures for books and music are similar (Caves 2000). The unpredictability of product appeal is not limited to cultural products. Gourville (2005) reports new product failure rates between 40 and 90 percent across many categories.

When the costs of bringing new products to market fall, society in effect can take more draws from an urn of potential new products. If the appeal of new products to consumers were perfectly predictable at the time of investment, then entry of additional products would be similar to adding more shelf space, virtual or otherwise, in a retail environment. The additional products would each have limited appeal and, in particular, lower appeal than the last currently entering product. But if appeal is unpredictable—and we will provide additional evidence that it is for music—then adding more products can have substantial benefits by delivering consumers products throughout the realized quality distribution. Because product appeal is also unpredictable in other industries, this idea may have broader applicability.

Technological change in the recorded music industry has allowed substantial growth in the number of new products, a tripling in new products between 2000 and 2008, leading us to explore how growth in the number

drid. All remaining errors are ours. Disclaimer: The views expressed are purely those of the authors and may not in any circumstances be regarded as stating an official position of the European Commission.

¹ By some estimates, the benefit consumers obtain from access to a long tail of additional varieties may be as high as \$1.03 billion per year for books alone in 2000 (Brynjolfsson, Hu, and Smith 2003).

of new products available affects welfare under our random long tail perspective in relation to the conventional view. We can measure the benefit as the difference between welfare with the new, enlarged choice set and a smaller choice set including a third of the recently entering products. Yet, the welfare impact of an entry cost reduction that triples the choice set depends heavily on which third of existing recent products would have entered without the cost reduction. This, in turn, depends on the predictability of quality at the time of investment. At one extreme, if product quality were perfectly predictable (the “perfect foresight” or PF case), then a reduction in the cost of entry from, say, T to T' would elicit entry of new products with expected—and realized—revenue between T and T' . The addition of these modest-appeal products to the choice set corresponds to the traditional long-tail benefits. The newly entering products would necessarily raise surplus available to consumers, but the benefit might be small since none of the new products would exceed the quality of the least attractive existing product. In the more realistic case in which quality is not entirely predictable (the “imperfect predictability” or IP case), benefits would be larger, as some new products would have high realized quality despite low expected revenue.

To quantify the benefits of new products made possible by digitization, we develop a simple illustrative equilibrium model of the recorded music industry that includes a structural demand model and a model of entry based on expected revenue. We use data on digital music track sales for 17 countries in 2011 to estimate a nested logit model of demand. The output of the model includes both parameter estimates and measures of the realized appeal of each product, which we term δ . We use the realized δ 's for the United States in 2011 to develop a forecasting model of expected quality, which we incorporate in our entry model. We infer fixed costs from the expected revenue of the last entering product. The model allows us to address the two questions that motivate the paper. First, what is the effect of the cost reductions associated with digitization—which have tripled the number of products brought to market in the United States—on consumer surplus and overall welfare? And our second, main question is, How do these benefits, which we term the long tail in production, relate to the conventional long tail in consumption?

Despite our detailed data on track sales, some features of our context place limits on the richness of the demand model we can estimate. Hence, our exercise is best viewed as an empirical illustration of the idea of the long tail in production rather than a richly specified demand estimation exercise. Even if we cannot precisely measure the size of the welfare benefit from new products, we can make stronger statements about its size in relation to the conventional long-tail benefits associated with the internet. We find that the size of the long tail in production relative to the conventional long tail is, perhaps surprisingly, quite insensitive to different

parameter estimates and demand modeling approaches. A tripling of the choice set according to expected quality adds nearly 20 times as much consumer surplus and more than 10 times as much overall welfare as a tripling of the choice set according to realized quality. That is, the long tail in production is almost 20 times as large as the traditional long tail. While insensitive to estimated demand parameters, the result does depend crucially on the predictability of quality, and we explore the sensitivity of our finding to different degrees of predictability. It is hard to know precisely how well industry participants can predict quality, but evidence presented here and elsewhere about the unpredictability of quality leads us to the conclusion that the random long tail is substantially larger than its conventional counterpart.

The paper proceeds in seven sections after the introduction. Section II presents descriptive facts about entry in the music industry, institutions for product discovery in the digital era, and a simple model illustrating the impact of unpredictability on the welfare effects of entry. Section III sets out an empirical structural model of the music market. Section IV presents the data that we will use in our estimation, while Section V presents our estimates of demand, expected revenue, and the fixed costs from the entry model. Next, we turn to counterfactual results in Section VI, including both estimates of the welfare impact of an enlarged choice set with imperfect predictability and our main object of interest, the size of this welfare gain in relation to the welfare impact of an enlarged choice set with perfect foresight. Section VII discusses the sensitivity of results to estimated parameters and forecasting approaches. Section VIII presents conclusions.

II. Background

A. *Industry Background*

This subsection provides background on three issues relevant to our exercise. First, we discuss technological change and the growth in new products. Second, we provide information on institutions for product discovery. Third, we describe existing evidence on the unpredictability of product quality.

Since 1999, recorded music revenue has fallen by 70 percent around the world. While industry participants—particularly the major record labels—have raised concerns that declining revenue would undermine investment incentives, the number of new products brought to market has risen rather than fallen as the cost of bringing new products to market has fallen substantially. As documented elsewhere, the costs of production, promotion, and distribution of new music have fallen sharply with digitization. These cost reductions are substantial enough to have enabled growth

in the number of new products despite the drastic decline in revenue; and the number of new recorded music products brought to market each year has risen since 1990 and more sharply since 2000 (Oberholzer-Gee and Strumpf 2010; Handke 2012; Waldfogel 2015; Aguiar and Waldfogel 2016a). According to Nielsen data, the number of new music products brought to market tripled between 2000 and 2008.²

The welfare that society derives from music equals the benefit to consumers, beyond what they pay, plus producer surplus, less costs of production and product discovery. With the substantial growth in new products, we would expect product discovery costs to rise. In particular, one might expect difficulty in consumer discovery of good products among the plethora of new offerings. Indeed, it is possible that consumers would fail to discover good products among the new releases, particularly among the new products released without much fanfare (e.g., little-known artists on independent labels). Under our imperfect predictability view of the world, we would expect some of the products with low *ex ante* promise to be highly valuable to consumers if they were discovered. If products with modest *ex ante* promise make up a growing share of the new music that becomes commercially successful, then we would infer that new products do not overwhelm the new product discovery institutions. And, indeed, in related research this is exactly what we find: products from independent labels, as well as products from new artists, make up growing and now substantial shares of the best-selling new recorded music (Waldfogel 2015; Aguiar and Waldfogel 2016a). In addition to providing evidence that consumers can discover new products, the growing success of the modest *ex ante* promise products also confirms the unpredictability of product success.

B. Related Literature

The study quantifies the benefit of a technological change that allows more entry of new products, given that new product appeal is unpredictable. Our question and approach are related to both the literature estimating the welfare effects of particular new products and the entry literature. Many studies evaluate the welfare impact of new products. A few prominent examples include Hausman and Leonard (2002) and Petrin (2002). The usual approach is to estimate demand in the presence of the new product and then to simulate welfare in the absence of the new product. We similarly do that, but we also model the entry process. That is, the comparative static that we evaluate is not simply about whether a particular new product—such as the minivan or breakfast cereal—exists, but rather about the cost of entry that would give rise to new products.

² See, e.g., <http://tinyurl.com/how-many-releases>.

Our paper is therefore closely related to the strand of the entry literature that incorporates demand modeling and therefore allows for explicit estimates of fixed costs (e.g., Berry and Waldfogel 1999). Usually, researchers postulate a model in which products (or firms) enter as long as their variable profit exceeds their fixed costs, and fixed costs are estimated from the expected revenue of marginal entrants. Observed entry configurations can then be viewed as Nash equilibria given the estimated fixed costs. Such models can be used to estimate welfare under, say, counterfactual fixed costs. Our exercise does this, adding the novel feature that product appeal is unpredictable at the time of entry.

Our exercise is also closely related to the literature on the “long-tail” benefits of the internet. Brynjolfsson et al. (2003) quantify the benefit of access to the full list of books at Amazon in contrast to, say, the 100,000 books locally available to a consumer. The literature takes the view—implicitly or explicitly—that digitization raises the variety available to people via an infinite shelf space mechanism rather than the new product mechanism that we explore.³

C. How Would Entry Cost Reduction Affect Welfare?

To fix ideas this section describes the intuition of our approach. Section III discusses the explicit model. When entry costs are T , then all products with expected revenue above T enter, while those with lower expected revenue do not; when the entry cost falls from T to T' , then more entry occurs. Having more products in the choice set raises welfare, but the size of the impact of additional products on welfare depends on the predictability of product quality at the time of investment. To see this, consider the following simple model of product entry with the possibility of quality unpredictability.

At the time of investment, an investor forms an estimate of a product's marketability as the true revenue y plus an error ν : $y' = y + \nu$.⁴ If the entry cost is T , then all products with expected revenue $y' > T$ enter. If the entry cost T falls to T' , then all products with $y' > T'$ enter. When product quality is perfectly predictable ($\nu = 0$), then a reduction in entry costs brings new products with expected and realized revenue—and therefore, we infer, product quality—between T and T' . In the more realistic case

³ Sinai and Waldfogel (2004) show that locally isolated consumers make greater use of the internet. Anderson (2006) popularized the idea of the long tail in a book asserting that the long list of products at the tail of the distribution are growing in importance relative to the small number of products at the head. Quan and Williams (2017) make the point that one will overstate the long-tail benefits of access to a large choice set if one overlooks the fact that offline assortments are tailored to local tastes.

⁴ This setup is reminiscent of Terviö (2009).

in which product quality is not perfectly predictable at the time of investment, the addition of products with expected revenue between T and T' elicits entry of products whose realized revenue might be anywhere in the distribution and can, of course, exceed T .

Our main concern in this paper is the evaluation of an entry cost reduction that tripled the number of new products. Given that digitization has already occurred, the welfare effect of digitization is the difference between the welfare associated with the current status quo choice set and the choice set including only a third as many new products. The major challenge to this exercise, however, is determining which third of recently added status quo products would have existed if digitization had not reduced entry costs. This, in turn, depends on the predictability of product quality.⁵

If investors had perfect foresight—and product quality were therefore completely predictable to investors at the time of entry—then when costs were high, only the products with the highest expected and realized quality would enter. Hence, the counterfactual high-entry cost choice set would be the top third of products according to realized quality. The comparison of the top third of products with the total choice set is analogous to the shelf space problem underlying the usual long-tail welfare calculation asking, for example, what benefit consumers derive from access to the top million books as opposed to the top 100,000. Under this usual approach, the benefit of additional products would be relatively small. At the other extreme, if quality were completely unpredictable to investors, then the counterfactual choice set associated with high entry costs would be a random sample of status quo products. Because the additional products would be as good, on average, as existing products, the additional products would add more to welfare than if investors had perfect foresight.

In the more plausible intermediate case of imperfect predictability, the effect of new products on welfare would fall between the two polar cases. This discussion demonstrates that the impact of cost reduction on product entry and resulting welfare—the long tail in production—depends crucially on the predictability of product quality to investors.

III. The Model

This section describes the components of our equilibrium model of the recorded music industry needed for measuring the welfare impact of the cost reduction: demand, quality prediction, and entry.

⁵ While we focus throughout the text on the welfare benefit of tripling the number of new products, we also explore the welfare consequence of different degrees of growth in the choice set. See Sec. VII.

A. Demand

Given our goal of developing an entry model incorporating expectations about product quality, we employ a model that allows us to easily infer product quality while also allowing for substitutability among products. To this end we employ a nested logit model, similar to that of Berry (1994) and Ferreira, Petrin, and Waldfogel (2013).

In each country, consumers choose whether to buy music and then choose among available songs. The choice sets of songs vary both across countries and over time. Define J_{ct} as the set of songs available in country c at time t , and index songs by j .⁶ Suppressing the time subscript, each consumer therefore decides in each month whether to download one song in the choice set $J_c = \{1, 2, 3, \dots, J_c\}$ or to consume the outside good (not purchasing a song). Specifically, every month every consumer i in country c chooses j from the $J_c + 1$ options that maximizes the conditional indirect utility function given by

$$\begin{aligned} u_{ij} &= x_{jc}\beta - \alpha p_{jc} + \xi_{jc} + \zeta_i + (1 - \sigma)\epsilon_{ij} \\ &= \delta_{jc} + \zeta_i + (1 - \sigma)\epsilon_{ij}, \end{aligned} \quad (1)$$

where δ_{jc} is therefore the mean utility of song j in country c . The vector x_{jc} includes song- as well as country-specific characteristics relevant to consumer interest in the song (and therefore also relevant to the song's prospect for success ex ante), p_{jc} is the price of song j in country c , and α is the marginal utility of money. The parameter ξ_{jc} is the unobserved (to the econometrician) quality of song j from the perspective of country c consumers and can differ across countries for the same song (song j , e.g., can have different quality to US vs. French consumers), and ϵ_{ij} is an independent taste shock. In contrast to a simple logit model, the nested logit allows for correlation in consumers' tastes for consuming digital music.⁷

Given the functional forms associated with nested logit, we can calculate the market share and revenue of each product for any set of product qualities δ_{jc} .⁸

⁶ Our data cover only digital singles, not albums. See Sec. IV for details on the data.

⁷ In the logit model the individual taste ϵ_{ij} is independent across both consumers and choices and the conditional indirect utility function is given by $u_{ij} = \delta_{jc} + \epsilon_{ij}$. This prevents the possibility that consumers have heterogeneous tastes, i.e., differ in their taste for consuming music. We provide additional detail in Aguiar and Waldfogel (2016b).

⁸ In the nested logit model we can calculate δ_{jc} as

$$\ln(S_j) - \ln(S_{0c}) - \sigma \ln\left(\frac{S_j}{1 - S_{0c}}\right),$$

where S_j represents the market share of song j in country c and S_{0c} is the market share of the outside good. See Sec. V for details of implementation. We also implement our estimates using a plain logit model to explore robustness of the results to the demand model specification. See Sec. VII below.

B. *Quality Prediction*

Our model of entry with unpredictable product quality requires us to have a measure of the appeal, or commercial success, that an investor would expect from releasing song j . The results from our demand estimation allow us to construct estimates of the mean utility of each song as well as their quality predictions. For each song j , and omitting country subscripts, δ_j reflects the appeal it generates for consumers based on its market share. The explanatory variables x_j included in the demand model, which describe consumers' demand for the product, are also relevant to quality prediction. This gives rise to our first forecasting approach: With the x_j variables included directly in the demand model, we can recover a quality prediction directly from the estimated demand model in a single step.⁹ We note here that some variables that will be helpful for prediction—namely, record label identities—are available only for the United States, so we will also employ a second, two-step approach, recovering δ_j from the demand model, and then regressing it on both x_j and the label variables.

C. *Supply and Fixed Costs*

Our measure of the welfare associated with an entry configuration, or set of products that enters, is the sum of consumer surplus and revenue less the number of products times the fixed cost per product. The demand model gives us consumer surplus and revenue for any entry configuration. In order to evaluate the welfare associated with a set of entering products, we need fixed costs and the ordered set of entering products, which our entry model delivers. While the imperfect prediction model is our central approach and the approach we view as realistic, we also develop approaches using perfect foresight and no predictability, both to illustrate the intuition of our approach and to compare our estimates of the long tail in production with estimates analogous to the long tail in consumption, reflected in the perfect foresight model.

1. Perfect Foresight

Under perfect foresight (PF), products enter in order of realized quality, or δ_j . The fixed cost under the status quo is the expected revenue of the last (N th) entering product.

To estimate the counterfactual PF fixed costs that give rise to one-third of recent status quo entry, we must calculate the expected revenue of the last product when only the $N/3$ best-selling products enter. To this end, define δ_j as the realized quality of product j , and define Δ_j as the set of

⁹ In general, $\delta_j = x_j\beta - \alpha p_j + \xi_j$. Because our data do not include prices, αp_j becomes part of the constant term in x_j in our empirical implementation. Our prediction of song j 's quality is therefore $\delta'_j = x_j\beta$. See Sec. V for details of implementation.

products $\{\delta_1, \dots, \delta_j\}$. Because products are imperfect substitutes, revenue to each product depends on the full set of products in the market. The expected revenue to product 1 entering alone depends on Δ_1 , and so on. That is, if $E[r_k]$ is the expected revenue of product k , then $E[r_k]$ is a function of the vector Δ_k .

If we order the products such that $\delta_k > \delta_{k+1}$, the products enter as long as $E[r_k(\Delta_k)] > T$, where T denotes fixed costs of entry. For example, given the nested logit structure, the expected and realized revenue to product 1 when it is alone is

$$r_1 = pMs_1 = pM \left[\frac{e^{\delta_1/(1-\sigma)}}{D_1^\sigma + D_1} \right], \quad (2)$$

where

$$D_1 = e^{\delta_1/(1-\sigma)}.$$

The price of the product is denoted by p , and M is market size.¹⁰

More generally, the revenue to product k (when it is the last entering product) is given by

$$r_k = pMs_k = pM \left[\frac{e^{\delta_k/(1-\sigma)}}{D_k^\sigma + D_k} \right], \quad (3)$$

where

$$D_k = \sum_{j=1}^k e^{\delta_j/(1-\sigma)}.$$

To estimate counterfactual fixed costs when $N/3$ products enter, we can infer that the fixed costs (T) equal the expected (and realized) revenue of the last entering product: $T \approx r_k$, with $k = N/3$.¹¹

2. No Predictability

At the opposite extreme from the perfect predictability model is a model with no predictability. While not a plausible depiction of reality, this model nevertheless provides a useful benchmark, describing a world in which,

¹⁰ In our empirical implementation, we define the market size as 12 times the country population. We also explore the sensitivity of our results to different market size definitions. See Sec. V.A.4.

¹¹ Our PF fixed-cost estimates require an important caveat (which applies to our IP estimates as well). We derive our estimates of fixed costs from the expected revenue of the marginal entering product. Hence, strictly speaking, our fixed cost is an estimate of the fixed cost for the marginal entrant. It seems likely that inframarginal entrants incur higher fixed costs. This means, further, that our estimate of the aggregated fixed costs incurred by all entrants, $N \cdot FC$, is a lower bound on the resources consumed by the fixed costs of entry. Underestimation of $N \cdot FC$ would lead to overestimates of welfare. We can, however, place an upper bound on fixed costs as well. Under free entry, entry could occur until profit opportunities have been dissipated. Hence, total revenue itself provides an upper-bound estimate of aggregate fixed costs ($N \cdot FC$). See Sec. V.A.3.

literally, “nobody knows anything.” With no predictability, all products are identical ex ante. Hence the expected revenue of any product depends only on the total number of products entering (k) and is the total revenue to those k products divided by k . That is,

$$E[r_k] = pME \left[\frac{D_k}{D_k^\sigma + D_k} \right], \tag{4}$$

where D_k is evaluated with a particular draw of k product qualities (δ_j), p is the price, and M is market size.

Hence, the no-prediction estimate of status quo fixed cost is the total observed revenue divided by the number of products. We estimate counterfactual fixed cost as the average revenue per product if $N/3$ products entered. To estimate this, we take draws of $N/3$ δ 's, and each draw generates an estimate of average revenue per product.

Under the no-predictability model, additional products add substantially to welfare by construction because the average quality of products does not decline with entry. The only reason that consumer surplus and the expected revenue per product decline with entry is that substitution is allowed for by the nested logit model's parameter σ .

3. Imperfect Prediction

The perfect foresight and no-prediction models present two extremes, both somewhat unrealistic. This leads us to the imperfect prediction case, in which investors have some ability to predict the appeal of songs at the time of investment. Our predicted δ 's (which we term δ') create an ordering of potential projects in descending order of ex ante (expected) promise: $\delta'_1 > \delta'_2 > \dots > \delta'_N$. In the imperfect prediction case, the analogue to the no-prediction random draw of k products is the top k products ordered by expected quality.

We calculate the expected revenue of the k th entrant as follows. Order songs by their ex ante promise (δ'). When the first $(k - 1)$ songs, ordered by their ex ante appeal, are in the market with their ex post appeal, the revenue to the k th entrant depends on its realized value. For a particular realization of $\delta_k = \delta'_k + \varepsilon$, the share of population consuming product k , via the nested logit formula, is

$$s_k(\varepsilon) = \frac{e^{(\delta'_k + \varepsilon)/(1-\sigma)}}{\left[\sum_{j=1}^{k-1} e^{\delta'_j/(1-\sigma)} + e^{(\delta'_k + \varepsilon)/(1-\sigma)} \right]^\sigma + \left[\sum_{j=1}^{k-1} e^{\delta'_j/(1-\sigma)} + e^{(\delta'_k + \varepsilon)/(1-\sigma)} \right]}. \tag{5}$$

Because of the nonlinearity of the share formula, we compute the expected market share by integration. The expected market share of the k th entrant is therefore given by

$$E[s_k] = \int s_k(\varepsilon)f(\varepsilon)d\varepsilon, \quad (6)$$

where f is the density of ε . In our empirical implementation, we will take f to be the empirical distribution of the residuals from our prediction model, $\varepsilon \equiv \delta - \delta'$. We will therefore compute the expected revenue of the k th entrant (when the first $k - 1$ songs ordered by their ex ante appeal have entered) as

$$E[r_k] = pME[s_k] = pM \left[\frac{1}{N} \sum_{n=1}^N s_k(\varepsilon_n) \right] = pM \left[\frac{1}{N} \sum_{n=1}^N \frac{e^{(\delta_k + \varepsilon_n)/(1-\sigma)}}{D_{kn}^\sigma + D_{kn}} \right], \quad (7)$$

where

$$D_{kn} = \sum_{j=1}^{k-1} e^{\delta_j/(1-\sigma)} + e^{(\delta_k + \varepsilon_n)/(1-\sigma)}$$

and N is the total number of products.

We estimate status quo fixed costs using the expected revenue of the last entrant, and we estimate counterfactual fixed cost as the expected revenue of the last product when the top $N/3$ products enter according to expected quality, or $k = N/3$.

IV. Data

Given our goals of estimating the welfare benefits of new music products, we would ideally observe all revenue generated by new music products. This would include sales of digital music, sales of physical products (e.g., compact disks), as well as live performance revenue. Our actual data, while very rich and detailed, include only a subset of the ideal. That is, the basic data for this study include annual sales of all digital singles in the United States, Canada, and 15 European countries, 2006–11, but our data contain no information on physical products or live performance revenue.¹² Our sample includes 3,984,227 distinct tracks from 75,235 distinct artists and, because a song can appear in multiple countries and years, 50,828,216 observations. Total digital track sales in the data are 628.3 million in 2006 and rise to 1,512.4 million in 2011. The sales data are drawn from Nielsen's SoundScan product, and readers are directed to Aguiar and Waldfogel (2016a, 2016b) for additional details.

¹² The data set initially includes the following 16 European countries: Austria, Belgium, Denmark, Finland, France, Germany, Ireland, Italy, Netherlands, Norway, Poland, Portugal, Spain, Sweden, Switzerland, and the United Kingdom. However, given that Poland enters the data in 2008 only, we decided to drop it from the analysis.

We use these underlying data to create two data sets that we use for demand estimations and quality predictions, respectively. While we have data for 2006–11, the main data set used for demand estimation covers 17 countries for 2011 and includes data on artists' country of origin, artists' age (measured as the number of years between a song's release year and the artist's earliest vintage release), as well as an artist genre designation.¹³ We obtained the genre data from <http://Allmusic.com>.¹⁴ We perform the quality prediction exercise and welfare calculations using only the subset of US data since these data also include the identity of labels releasing each song. Our revenue data cover digital track sales, not the total revenue that artists earn from creating music. The track sales are a subset of total recorded music sales. Our US digital track sales total \$1.313 billion for 2011, while the Recording Industry Association of America (RIAA) reports total recorded music sales of \$7.008 billion.¹⁵ Hence, to make them reflective of US recorded music sales, we scale up our estimates by 5.34.¹⁶ We discuss this further in Section V.C below.

Other variables we employ in the study include population, GDP per capita, and the urban share of the total population. These are drawn from the World Bank Open Data (<http://data.worldbank.org/>). We also use measures of the digital share of music expenditure in each country and year, which we take from the *Recording Industry in Numbers* 2013 publication from the International Federation of the Phonographic Industry (<http://www.ifpi.org>). Descriptives on the variables used in the estimation can be found in Aguiar and Waldfogel (2016b).

V. Empirical Implementation

A. Demand Model

Following equation (1) and normalizing the utility of the outside good δ_{0c} to zero, the market shares for all $j \in J_c$ are given by

¹³ We rely on the 2011 data for our demand estimation because our identification strategies are cross-sectional and because the 2011 data are the most recent. See Sec. V.A.4.

¹⁴ We sought matches for each of the 75,235 sample artists from Allmusic.com. We obtained matches for 61,073 artists, accounting for 93.2 percent of the sales in the data with origin matches. The artists are classified into 36 distinct genres, which we aggregate to five broad genres: electronic, jazz, pop/rock, rap/R&B, and other.

¹⁵ The RIAA reports sales of 1,306.2 million digital tracks, generating \$1,492.7 in revenue, or \$1.14 per track. Our data contain 1.149 billion US track sales. At \$1.14 per track, our data cover \$1.313 billion in track sales. See RIAA, 2011 Year-End Shipment Statistics (<https://www.musicamemoria.com/docs/riaa-2011-year-end-shipment.pdf>).

¹⁶ Artists also derive revenue from live performances as well as recorded music. In 2011, live performance revenue was \$4.35 billion. See 2011 Pollstar Year End Business Analysis, available at <http://www.pollstarpro.com/files/Charts2011/2011BusinessAnalysis.pdf>. To the likely extent that the creation of new music also brings the opportunity to generate some live performance revenue, measures of expected revenue based only on recorded music sales would understate the true expected revenue.

$$S_{jc} = \frac{e^{\delta_{jc}/(1-\sigma)}}{D_{jc}^\sigma + D_{jc}}, \quad \text{with} \quad D_{jc} = \sum_{j \in J} e^{\delta_{jc}/(1-\sigma)}.$$

Inverting out δ_{jc} from observed market shares as in Berry (1994) yields

$$\begin{aligned} \ln(S_{jc}) - \ln(S_{0c}) &= \delta_{jc} + \sigma \ln\left(\frac{S_{jc}}{1 - S_{0c}}\right) \\ &= x_{jc}\beta - \alpha p_{jc} + \sigma \ln\left(\frac{S_{jc}}{1 - S_{0c}}\right) + \xi_{jc}, \end{aligned} \tag{8}$$

so that an estimate of β , α , and σ can be obtained from a linear regression of differences in log market shares on product characteristics, prices, and the log of within-group share. The vector x_{jc} includes three kinds of variables. First, we include country-level variables: the 2011 digital share of music expenditure in the country, GDP per capita, the shares of population that are urban, fixed broadband subscribers, mobile cellular subscribers, and internet users. Second, we include artist-level variables: genre, artist's age, country of origin, whether the artist is new with this release, and the artist's sales in previous years. Third, we include terms in the age of the song.

As the determinants of appeal in the demand model, the variables x_{jc} are also relevant to quality prediction. Hence, we can effectively estimate the demand and prediction models in one step. The one-step approach faces a practical obstacle, however, as one important set of variables (the identity of labels releasing each song) is available only for the United States and not for other countries and hence cannot be used directly in demand estimation. Our estimate of the relative welfare gain from additional products will be sensitive to the explanatory power of the prediction model. Hence, we will undertake estimation in two ways. First, we do one-step estimation, including all of the available song and artist characteristics that are predictive of song success directly in (8). Second, we take a two-step approach, deriving δ_j from our estimate of σ from (8) and then regressing δ_j for new US songs in a second step on all relevant predictors, including both x_j and the label variables that are available only for the United States.

1. Identification of σ

The substitution parameter σ plays an important role in showing the benefits of additional products to consumers. It is helpful to note that the demand model (8) is a regression of $\ln(S_{jc}) - \ln(S_{0c})$ on, among other variables, $\ln(S_{jc}) - \ln(1 - S_{0c})$. Before even considering the possible endogeneity of the independent variable, it is worth observing that it is a regression of a function of S_{jc} on a related function of S_{jc} . Hence, our first

approach, ordinary least squares (OLS), could produce an upwardly biased estimate of σ (indicating close substitutability of products) for mechanical reasons. Instrumental variables (IV) can address this problem as well as the possible endogeneity of $\ln(S_{jc}) - \ln(1 - S_{0c})$.

We explore two broad IV strategies for identifying σ . The endogenous variable of interest on the right-hand side of the demand equation is $\ln(S_{jc}) - \ln(1 - S_{0c})$. We can get some intuition about identification of σ from noting that in a symmetric model—if all inside products had equal market shares— $S_{jc}/(1 - S_{0c})$ would equal $1/N$, where N is the number of products. In general, the market share of an individual product j is a function of the number of remaining products ($N - 1$), as the product must compete with the other $N - 1$ products. Just as the market share of a product j depends on the number of products entering, it also depends on the natural determinant of the number of entering products, the size of the market. Hence, one can imagine using either N or measures of market size, such as population, as instruments for $\ln(S_{jc}) - \ln(1 - S_{0c})$.

Two points about using the number of products as an instrument are in order. First, this is the simplest version of the IV approach used in Berry, Levinsohn, and Pakes (1995) and described in Nevo (2000), which in general entails using functions of the other products in the choice set as instruments. In this case, the function is simply the sum of the products, or N . Second, one can be concerned that N is itself endogenous. If markets with more entry have elevated unobserved taste for recorded music, then this instrument will lead to estimates that overstate the market expansion arising from entry (and therefore understating the size of the demand parameter σ).

On the other hand, if market size affects entry conditions but is not directly related to preferences, then instrumenting $\ln(S_{jc}) - \ln(1 - S_{0c})$ with market size will avoid this overstatement of market expanding effects of entry. Under the European Union restrictions on cross-country digital trade, the decision to make a song available must be made on a country-by-country basis, and we therefore expect to see larger N in larger markets, as the product entry decision is undertaken for more songs in larger markets.¹⁷

Simple figures illustrate our two basic IV approaches to identification of the demand model. The left panel of figure 1 depicts the relationship between market size, as measured by the log of population, and the number of songs available in each country in 2011. The relationship is clearly positive. We take this as evidence of variation in the choice sets that is driven by the size of the market, as opposed to the tastes of individuals.

¹⁷ This is the approach employed in Berry and Waldfogel (1999) and Gentzkow and Shapiro (2010) and discussed in Berry and Haile (2015) and Berry and Waldfogel (2016).




FIG. 1.—Identification of the demand model, 2011. The left panel shows the number of songs and market size. The right panel shows the share consuming and number of songs.

The right-hand panel of the figure shows the relationship between the log of the number of products and per capita consumption; this too is positive, but as discussed above, one can be concerned that the number of products entering is endogenous to the level of demand for music.

While the simplest version of the Berry et al. approach uses the number of products as instruments, more complicated variants involve functions of the characteristics of products. Here, for our approach, we can use the sum of the following product characteristics: age, genre, and country of origin. In what follows, we report four groups of estimates of the demand model (8): OLS, IV using market size, IV using the number of products, and IV using functions of characteristics.

2. Price Coefficient

While we need an estimate of the price coefficient to translate the utility gain from additional products into a dollar value, the parameter plays no role in our calculation of the value of the random long tail relative to the conventional one. To estimate α , we would ideally observe exogenous price variation across songs that would allow us to econometrically identify the price coefficient α . This approach is infeasible because we do not observe song-level prices. We do, however, observe the average price, allowing us to infer the α parameter from a first-order condition on pricing.

Because the price is constant, the term αp_{jc} in (8) simply becomes part of the constant term in the estimating equation

$$\ln(S_{jc}) - \ln(S_0) = x_{jc}\beta + \sigma \ln\left(\frac{S_{jc}}{1 - S_{0c}}\right) + \xi_{jc}. \quad (9)$$

Using σ we can calculate the country-specific mean utility of each song δ_{jc} :

$$\delta_{jc} = \ln(S_{jc}) - \ln(S_{0c}) - \sigma \ln\left(\frac{S_{jc}}{1 - S_{0c}}\right). \quad (10)$$

We can infer α from a condition on the music demand elasticity. Assuming that songs are sold by a profit-maximizing monopolist facing zero marginal cost, the price level would be set such that the demand for songs is unit elastic.¹⁸ Given that the elasticity of demand for music in our model is given by

¹⁸ Note that this way of inferring α is not uncommon among practitioners. As noted by Bøjrnstedt and Verboven (2013), one may want to verify whether elasticities are consistent with external industry information as opposed to relying too heavily on econometric estimates. While our motivation is driven by lack of data on product prices, we basically follow the same type of approach.

$$\eta = \alpha p \left(1 - \frac{D_j}{D_j + D_j^\sigma} \right),$$

we can infer the price parameter under the assumption of unit-elastic pricing as

$$\alpha = \frac{1}{p} \frac{D_j + D_j^\sigma}{D_j^\sigma}.$$

In reality, it is likely that major sellers of digital music (e.g., Apple) price songs below the static profit maximization level to stimulate demand for complementary hardware (Shiller and Waldfoegel 2011; Danaher et al. 2014). If so, the estimate of α is an upper bound, and our resulting estimates of consumer surplus will be a lower bound.¹⁹

3. Consumer Surplus, Revenue, and Welfare Measures

Given our estimates of σ and α , we can calculate the mean utility of each song, and given these estimates of $\delta_{j\sigma}$, we can calculate the consumer surplus (*CS*) and revenue (*Rev*). These, in turn, allow us to calculate two kinds of welfare measures, *CS* and overall welfare, $W = CS + \text{Rev} - N \cdot FC$, where N is the number of products and FC is the fixed cost per product. Note that if entry costs equaled revenue, then welfare would simply equal consumer surplus. In what follows, we calculate the change in welfare both assuming that fixed costs are determined by the marginal entrant and assuming that fixed costs equal revenue, in which case $\Delta W = \Delta CS$. Use of consumer surplus as a welfare measure is also consistent with the literature in this area (e.g., Brynjolfsson et al. 2003), which focuses entirely on *CS*. Consumer surplus is given by

$$CS = \frac{M}{\alpha} \ln(D_j^{1-\sigma} + 1),$$

and revenue is given by²⁰

¹⁹ If demand is inelastic, then

$$p\alpha \left(1 - \frac{D_j}{D_j + D_j^\sigma} \right) < 1$$

and

$$\alpha < \frac{1}{p} \frac{D_j + D_j^\sigma}{D_j^\sigma}.$$

²⁰ The results from our estimations allow us to calculate *CS* and revenue for each country in each year. However, we omit the country and time subscripts since we perform our counterfactual exercise on US data in 2011 only.

$$\text{Rev} = p_j M \left(\frac{D_j}{D_j^\sigma + D_j} \right).$$

Our objects of interest are the absolute change in welfare with the new products and, especially, the change in welfare under our IP approach, relative to the standard PF long tail. Given our setup,

$$\Delta CS_{IP} = \frac{M}{\alpha} [\ln(D_j^{1-\sigma} + 1) - \ln(D_{j_0^{IP}}^{1-\sigma} + 1)],$$

where J is the full status quo choice set, and J_0^{IP} is the set of products that would have existed without a cost reduction under IP. Note that this depends on α , σ , and our predictions of which products enter the counterfactual IP choice set: $\Delta W_{IP} = \Delta CS_{IP} + \Delta \text{Rev}_{IP} - N_0 \cdot FC_0 - N_{IP} \cdot FC_{IP}$, where ΔRev_{IP} is the status quo revenue less the revenue that the top third of products in expected revenue would generate, and N_0 and FC_0 are the number of products and the fixed costs per product in the status quo, respectively.

Our second and main object of interest is the welfare change ratio:

$$\frac{\Delta CS_{IP}}{\Delta CS_{PF}} = \frac{\frac{M}{\alpha} [\ln(D_j^{1-\sigma} + 1) - \ln(D_{j_0^{IP}}^{1-\sigma} + 1)]}{\frac{M}{\alpha} [\ln(D_j^{1-\sigma} + 1) - \ln(D_{j_0^{PF}}^{1-\sigma} + 1)]} = \frac{\ln\left(\frac{D_j^{1-\sigma} + 1}{D_{j_0^{IP}}^{1-\sigma} + 1}\right)}{\ln\left(\frac{D_j^{1-\sigma} + 1}{D_{j_0^{PF}}^{1-\sigma} + 1}\right)}, \quad (11)$$

and, analogously, $\Delta W_{IP}/\Delta W_{PF}$. These ratios depend on σ and the products predicted to enter the choice set. Notice that while $\Delta W_{IP}/\Delta W_{PF}$ depends on α , the ratio $\Delta CS_{IP}/\Delta CS_{PF}$ does not. Moreover, although σ enters (11), so that the welfare change ratios formally depend on σ , as it turns out below, the ratios are empirically almost invariant to σ . Hence, the welfare change ratio turns out not to be sensitive to estimates of the demand parameters α and σ (although, again, it will depend on the predictability of product quality).

4. Results

Although we have data on 2006–11, our main estimates rely only on 2011 data, for two reasons. First, our identification strategies are cross-sectional. Second, the 2011 data are the most recent, and by 2011 digital music has been more widely adopted. Table 1 reports first-stage estimates using each of the three IV approaches. In all cases the dependent variable is the log of the inside share $\ln(S_{ic}) - \ln(1 - S_{oc})$, and we cluster standard errors in all specifications at the country level. Column 1 uses the log of population as an instrument. As figure 1 showed, market size works, in the sense that

TABLE 1
DEMAND MODEL: FIRST STAGE

	(1)	(2)	(3)
ln(population)	-1.173*** (.13)		
ln(number of songs)		-2.289*** (.12)	
Sum of songs' ages			-.013*** (.00)
Number of songs from Germany			.158*** (.01)
Number of songs from France			-.166*** (.01)
Number of songs from Spain			-.868*** (.03)
Number of songs from United Kingdom			.154*** (.01)
Number of songs from United States			-.554*** (.01)
Number of electronic songs			-2.157*** (.06)
Number of jazz songs			.497*** (.03)
Number of pop/rock songs			.237*** (.01)
Number of rap/R&B songs			.607*** (.03)
Number of other songs			.689*** (.02)
Share of digital sales	-2.498** (1.16)	-2.061*** (.49)	1.072*** (.15)
GDP per capita	-23.154 (15.28)	9.072*** (2.48)	-23.296*** (.77)
Urban population	-.009 (.02)	.007 (.01)	.036*** (.00)
Age of the song	-16.818*** (3.18)	-22.334*** (1.74)	-23.698*** (1.68)
(Age of the song) ²	311.673*** (45.28)	368.587*** (30.39)	386.088*** (29.80)
Genre fixed effects	✓	✓	✓
Origin fixed effects	✓	✓	✓
F-statistic excluded instruments	76.488	385.021	7,588.996
p-value	.000	.000	.000
Observations	10,800,378	10,800,378	10,800,378

NOTE.—All specifications use 2011 data and include variables measuring artists' past sales, artists' age and its square, an indicator for new artists, and time since last release. Standard errors are clustered on country level and are in parentheses.

* Significant at the 10 percent level.

** Significant at the 5 percent level.

*** Significant at the 1 percent level.

the number of available products is greater in larger markets: the larger the market, the smaller the average inside share. The *F*-statistic for the instrument is 76.5. When we use the number of products as an instrument, the instrument also works, in the sense of bearing a strong relationship with the endogenous inside share. Markets with more products have, on average, smaller inside shares; and the *F*-statistic is 385.0. Using the full Berry et al.–style instruments, the sums of the ages of products, and the numbers of products by genre and origin, the *F*-statistic is 7,588.

Table 2 reports estimates of the demand model using OLS as well as the three IV approaches laid out above. Our demand specifications include all of the country, song, and artist-level variables. Note that we do not report all of the prediction coefficients in table 2 but defer their discussion until table 3 and Section V.B. The first point to observe is that OLS gives rise to our highest estimate of σ , 0.786, possibly for the mechanical reason that we are regressing a function of S_{j_c} on another function of S_{j_c} .

Estimation using the market size instrument, in column 2, gives a σ estimate of 0.751. Estimates using the simplest Berry et al. approach—that

TABLE 2
DEMAND MODEL

	ORDINARY LEAST	INSTRUMENTAL VARIABLES		
	SQUARES	(2)	(3)	(4)
	(1)			
$\ln[s_j/(1 - s_0)]$.786*** (.07)	.751*** (.10)	.533*** (.09)	.511*** (.08)
Share of digital sales	3.159*** (1.08)	2.889** (1.30)	1.205*** (1.07)	1.032*** (1.07)
GDP per capita	19.471 (11.97)	20.745* (11.73)	28.716*** (7.51)	29.534*** (7.33)
Urban population	.007 (.01)	.006 (.01)	.003 (.02)	.003 (.02)
Age of the song	-12.787*** (4.14)	-12.626*** (3.79)	-11.614*** (2.89)	-11.510*** (2.81)
(Age of the song) ²	169.008*** (57.40)	171.401*** (54.32)	186.377*** (43.54)	187.914*** (43.08)
Genre fixed effects	✓	✓	✓	✓
Origin fixed effects	✓	✓	✓	✓
Instruments	...	$\ln(\text{Pop})$	$\ln(N)$	Sums of age, origin, genre
<i>R</i> ²	.872	.871	.806	.794
Observations	10,800,378	10,800,378	10,800,378	10,800,378

NOTE.—All specifications use 2011 data and include variables measuring artists’ past sales, artists’ age and its square, an indicator for new artists, and time since last release. Columns 2, 3, and 4 use $\ln(\text{population})$, $\ln(\text{number of products})$, and Berry et al.–style instruments, respectively. Standard errors are clustered on country level and are in parentheses.

- * Significant at the 10 percent level.
- ** Significant at the 5 percent level.
- *** Significant at the 1 percent level.

is, using $\ln(N)$ as an instrument, in column 3—give 0.533. The estimates in column 4 using all of the Berry et al. instruments give a similar σ estimate of 0.511. Hence, the Berry et al. approach indicates a larger market expansion effect, that additional products add a great deal to consumption. This is consistent with the intuition about endogeneity of the number of products.

We have a range of demand estimates before us, and we need to choose an estimate for carrying out the simulations. OLS has the disadvantage of potential endogeneity, as well as a mechanical relationship between the dependent and independent variables of interest. Of the IV approaches, the market size instrument is most appealing to us. The use of the resulting σ estimate (0.751) has the justification that it is conservative, relative to the other IV approaches, in that the larger σ estimate will give rise to smaller absolute welfare benefits of digitization. Two other points are important. First, it bears continued emphasis that our relative welfare measure will turn out to be insensitive to σ . Second, we will calculate all results of interest for a range of σ values extending beyond the plausible range in table 2.²¹

B. Quality Prediction

While we estimate the demand model on data for 17 countries, we perform our counterfactual exercises on only US data for 2011. In our counterfactual calculations we treat only the vintage 2011 products as endogenous. That is, we treat the pre-2011 products available in 2011 as exogenously available and omit the bottom two-thirds of vintage 2011 products (according to their expected quality) in our counterfactual choice set. These simulations can be interpreted to represent a cost reduction that occurred starting in 2011.

Implementing our simulation requires quality predictions, and the first quality prediction comes straight from the demand model. Recall, from the demand model, that $\delta_j = x_j\beta + \xi_j$, so we have estimates of β from the demand model above. Column 1 of table 3 reports the estimates of the coefficients on the prediction variables from the one-step estimation approach. Because we are interested only in predicting the quality of 2011 songs, the demand model includes separate coefficients for the vintage 2011 US products and the other products. In table 3 we report only the coefficients on the interactions of vintage 2011 US dummies with the variables of interest. We see, for example, that new releases from artists with greater recent sales tend to have higher realized quality and that sales are

²¹ In Aguiar and Waldfogel (2016b) we explore a large number of alternative specifications, with similar results.

TABLE 3
FORECASTING MODEL

	(1)	(2)	(3)	(4)
ln(sales in $t - 1$)	141.575*** (53.85)	91.731*** (.96)	69.240*** (.89)	65.174*** (.78)
ln(sales in $t - 2$)	-5.858** (2.81)	-1.420 (1.13)	-4.633*** (1.03)	-12.153*** (.89)
ln(sales in $t - 3$)	-15.727*** (6.09)	-13.757*** (1.13)	-10.377*** (1.04)	-3.885*** (.86)
ln(sales in $t - 4$)	-7.860** (3.06)	-4.769*** (1.11)	-2.631** (1.02)	.040 (.66)
ln(sales in $t - 5$)	-8.509* (4.58)	-.103 (.84)	-4.858*** (.77)	
Years since last release	40.745*** (15.81)	14.557*** (.77)	9.757*** (.70)	8.609*** (.56)
Artist's age	-2.041 (5.98)	-14.678*** (.41)	-11.257*** (.38)	-12.617*** (.34)
(Artist's age) ²	-.056 (.08)	.161*** (.01)	.129*** (.01)	.152*** (.01)
New artist	1.254*** (.48)	.834*** (.01)	.614*** (.01)	.453*** (.01)
Genre fixed effects	✓	✓	✓	✓
Origin fixed effects	✓	✓	✓	✓
Label fixed effects	X	X	✓	✓
R^2238	.411	.403
Prediction R^2	.196	.238	.411	.323
Observations	10,800,378	134,241	134,241	156,411

NOTE.—Specification 1 reports the estimates of the coefficients on the prediction variables from the one-step estimation approach using 2011 data. Column 1 reports the coefficients on the interactions of vintage 2011 US dummies with the variables of interest. Columns 2 and 3 use 2011 data and songs from vintage 2011. Column 4 uses 2010 data and songs from vintage 2010. The predicted δ 's are constructed for the vintage 2011 songs in all specifications. The prediction R^2 is computed as the square of the correlation between the realized δ 's in 2011 and their prediction. Standard errors are clustered on country level and are in parentheses.

* Significant at the 10 percent level.

** Significant at the 5 percent level.

*** Significant at the 1 percent level.

lower for older artists. How much of the variation in realized quality does this prediction model explain? We are interested in the model's explanatory power only for vintage 2011 US products. Hence, we calculate the model's predicted quality for those products, and we calculate the relevant R^2 as $\text{corr}(x_i\hat{\beta}, \delta_i)^2$ on the vintage 2011 products. We are able to explain 19.8 percent of the variation in realized quality for vintage 2011 US songs with the one-step estimates.

The one-step estimates have the advantage of being derived from simultaneous estimation, but they have the disadvantage that we cannot include variables that are not available for all countries. As a result, it is important to know whether a two-step approach, which would allow us to

include the variables available only for the United States, yields similar results. To explore this, we derive an estimate of δ_j from our baseline demand estimation and then regress it directly on the x_j variables for just the vintage 2011 US observations.²² Column 2 reports results, and the coefficients are similar to those in column 1. Moreover, the R^2 of this regression is also similar to the implied R^2 for column 1 albeit somewhat higher, at .238 rather than .196. This provides some indication that the two-step approach will not produce misleading results.²³

Column 3 adds the label identifiers to the specification in column 2. The data contain 13,507 different labels. Artists tend to match with different labels according to expected quality, with the “major” labels releasing artists with substantial commercial appeal and the independents releasing artists with more modest prospects. There is, moreover, a range of independent labels from labels such as Merge and 4AD handling well-known “indie” artists to more obscure labels. Hence, label dummies should be correlated with predictors of success that labels can observe but the econometrician cannot.²⁴

There is an important sense in which all of the prediction models in columns 1–3 of table 3 overstate predictability. These models use 2011 data to predict the success of 2011 releases and may suffer from overfitting. The realized qualities of the 2011 releases are not known at the time of investment, and to mimic the decision problem of investors should use only information available prior to the realizations of the 2011 vintage releases’ success for prediction. To this end we can instead estimate the demand and prediction models on 2010 data and then use the resulting coefficients along with 2011 X ’s to form predictions. We calculate the prediction R^2 using the 2010 parameters. This R^2 is then $\text{corr}(x_{j,2011}\hat{\beta}_{2010}, \delta_{j,2011})^2$. Using the 2010 forecast reduces the R^2 from .411 in column 3 to .323, and this is the baseline estimation approach we will use in the paper; but we will explore the sensitivity of the results to predictability extensively in Section VII.C.

We make one final observation. Our prediction model tells which vintage 2011 products would not have been available to consumers in 2011 without cost reductions following from digitization. In reality, cost reduction—and the growth of new releases—predates 2011, so that the full

²² That is, we take our baseline estimate of σ (0.751) and then calculate δ_j using expression (10) above.

²³ Even more relevant for us, both prediction equations give rise to roughly equal values of the ratio in (11).

²⁴ To the extent that labels have already formed a prediction of the artist’s appeal when signing them, including label fixed effects in the forecasting model will arguably lead to more conservative results. The R^2 rises fairly substantially with the inclusion of label dummies, to .411.

benefit of digitization that US consumers experience during 2011 exceeds the benefits associated with the additional vintage 2011 products.

C. Fixed Costs

Our estimates of fixed costs are based on estimates of the expected revenue of the last entering US product. That is, we calculate perfect foresight status quo fixed costs as the expected (and realized) revenue of the lowest-appeal vintage 2011 product. Because the lowest revenue observed in the US digital song data for a vintage 2011 song is \$1.14, the resulting status quo fixed-cost estimate under perfect foresight is \$1.14. Scaling this up to the total year 2011 US recorded music revenue (multiplying by 5.34) yields a fixed cost of \$6.09 (see table 4). The analogous perfect foresight counterfactual fixed cost is estimated as the expected revenue of the last entering vintage 2011 product when all pre-2011 products are in the choice set while only the top third of vintage 2011 products (by realized quality) enter. We estimate this to be \$133.97.

We estimate status quo no-predictability fixed costs by calculating the average revenue to each of the vintage 2011 products when they are available alongside the earlier, exogenous products (from vintages prior to 2011). We estimate these as \$9,467.89. For the counterfactual no-predictability fixed costs, we randomly remove two-thirds of the vintage 2011 songs and then calculate the average revenue per 2011 song when, again, they are sold alongside all of the pre-2011 songs. We repeat this random exercise 5,000 times, resulting in a counterfactual fixed-costs estimate of \$10,521.81.

We calculate the imperfect predictability status quo fixed costs by ordering the 2011 products by expected quality. We then seek an estimate of the expected revenue of the last entering vintage 2011 product when it is available alongside both the preceding vintage 2011 product and all

TABLE 4
FIXED COSTS OF ENTRY

Regime	Perfect Foresight	Imperfect Predictability	No Predictability
Counterfactual	133.97	1,792.23	10,521.81
Status quo	6.09	18.97	9,467.89

NOTE.—Fixed costs are estimated as the expected US digital single revenue of the last entering product, scaled up to the size of the entire US recorded music market in 2011. Status quo refers to the set of products available in the United States in 2011, while counterfactual models the choice set if digitization had not occurred, referring to simulations in which the bottom two-thirds of vintage 2011 products, by expected revenue, are removed from the choice set. Under perfect foresight, products are ordered by realized revenue. Under our main model, imperfect predictability, products are ordered by expected revenue. With the no-prediction model, products are ordered randomly (so that the counterfactual choice set has one-third of actual vintage 2011 products, chosen at random). All figures are in \$US 2011.

of the pre-2011 products. Using equation (7), we estimate the status quo fixed cost as the expected revenue of the last entering product. We obtain an estimate of \$18.97. We similarly estimate the counterfactual imperfect predictability fixed costs as the expected revenue of the $k = (N/3)$ rd entering product, obtaining an estimate of \$1,792.23.

As is customary in the empirical entry literature, our fixed-cost estimates are derived from a cross section of revenue data. The fixed costs derived from year 2011 expected revenue of new vintage 2011 songs reflect only expected first-year song revenue. If first-year revenue is proportional to lifetime revenue, then our fixed-cost estimates will be proportional to the true underlying fixed costs. Moreover, the fixed-cost estimates derived from first-year revenue bear the same relationship to total fixed costs that our observed first-year revenue bears to total revenue. Hence, revenue and cost estimates are consistent with one another, for example, for the purpose of entry counterfactuals involving different fixed-cost levels.²⁵

While the status quo and counterfactual fixed-costs estimates are mainly inputs into our welfare calculations, they are also of some direct interest as answers to the question, How much must fixed costs have fallen to generate a tripling of entry? The answer, under imperfect predictability, is roughly a factor of 100, from about \$1,800 to \$19 in table 4.

VI. Simulations: Effect of Tripling the Number of Songs on Welfare

Table 5 reports baseline estimates of both the absolute changes in welfare measures (ΔCS and ΔW) and our main objects of interest, the ratios $\Delta CS_{IP}/\Delta CS_{PF}$ and $\Delta W_{IP}/\Delta W_{PF}$. Using our imperfect predictability approach, the additional vintage 2011 songs in the 2011 choice set raise CS by \$10.09 million, and given the implied reduction in entry costs, W rises by \$71.72 million. Recall that these are inflated to reflect total US recorded music sales.

The PF welfare benefits of additional entry, corresponding to the traditional long tail, are far smaller than the IP benefits. Consumer surplus for 2011 rises by \$0.51 million with a tripling in the number of new 2011 products, while W rises by \$6.20 million. These absolute changes in welfare, in addition to reflecting any uncertainty we have about the true value of the substitution parameter, are also small. It is important to note that we calculate the absolute change in consumer surplus by removing the least-promising 90,000 vintage 2011 products from a 2011 choice set that

²⁵ One might want to adjust our fixed-cost estimates to compare them to outside estimates of the cost of bringing new music to market. We do so by considering additional sources of revenue that released products generate. In particular, we consider the revenue obtained from the full life of the song as well as the revenue from live performance. See Aguiar and Waldfogel (2016b) for details on the calculations of these fixed-cost adjustments.

TABLE 5
COUNTERFACTUAL RESULTS

Regime	ΔCS	ΔCS Ratio	ΔRev	ΔRev Ratio	ΔTC	ΔTC Ratio	ΔW	ΔW Ratio
Perfect foresight	.51	1	.51	1	-5.18	1	6.20	1
Imperfect predictability	10.09	19.82	10.09	19.82	-51.55	9.96	71.72	11.57
No predictability	152.42	299.48	153.16	300.93	800.16	-154.54	-494.58	-79.82

NOTE.— ΔCS is the change in CS from the tripling of the vintage 2011 products made possible by digitization. The three regimes differ by which products are in the counterfactual (no-digitization) choice set. Perfect foresight adds products with the lowest realized quality, while imperfect predictability adds products with the lowest expected quality. The no-predictability regime adds products that are as good, on average, as the products that would be available without digitization. ΔCS Ratio reports ΔCS relative to the perfect foresight estimate that corresponds to the traditional long tail; ΔRev , ΔTC , ΔW , and the corresponding ratios are defined analogously. TC is the fixed cost per product times the number of entering products.

continues to include over 2 million products. Given the substitutability across products, as well as the fact that pre-2011 products collectively account for a large share of sales, we would not expect the absolute increase in welfare to be large. Again, the absolute change in welfare under imperfect predictability interests us mainly in relation to its analogue under perfect foresight. And using our baseline model, our long tail in production produces a ΔCS_{ip} benefit that is 19.82 times larger than traditional perfect foresight benefit ΔCS_{pf} . Our overall welfare benefit ΔW_{ip} is 11.57 times larger. This is our main finding.

Although the absolute size of the welfare gain is not our main focus, a few notes are in order. We estimate that consumers experienced a \$10.09 million benefit in 2011 from the vintage 2011 products made possible by digitization. During 2011, US consumers also enjoyed additional new products released in 2010, 2009, 2008, and so on, back to 2000 if one were to mark the onset of digitization following Napster. A rough estimate based on the growth in the number of new products since 1999 and the shares of these vintages in year 2011 sales suggests that the role of new, digitization-enabled products in the 2011 choice set is 4.31 times as large as the new 2011 products alone.²⁶ Hence, the full year 2011 benefit of new products is roughly four times the benefit arising from just the new (vintage 2011) products. This is \$43.49 million for the United States in 2011, a year in which total recorded music sales were \$7 billion.

Our estimates of the absolute size of the welfare benefits from new products appear small in comparison with existing long-tail estimates. Brynjolfsson et al. (2003) estimate that access to all book titles at Amazon, rather than just the top 100,000 titles, delivered \$1 billion in additional consumer

²⁶ See the online appendix for details on these calculations.

surplus to US consumers in 2000. Their measurement approach corresponds to what we term perfect foresight but applied to all vintages rather than just the 2011 vintage. Our basic PF approach counterfactually removes the lowest-demand two-thirds of products released in 2011. We can produce an estimate more closely resembling Brynjolfsson et al.'s approach by discarding all but the 100,000 most popular products among the full 2.2 million products available in 2011 regardless of vintage. The loss in CS from eliminating all but the top 100,000 products is \$86.4 million. This figure remains smaller than the corresponding measure for books, largely because books have far lower sales concentration. Brynjolfsson et al. report that books outside the top 100,000 titles accounted for about 40 percent of book sales in 2000. In our music data, tracks outside the top 100,000 account for under 5 percent of sales. Hence, we expect our estimates of conventional long-tail benefits (the benefits arising from access to products outside, say, the top 100,000) to be much smaller than a corresponding estimate for books.

VII. Robustness

Our estimates of the absolute changes in welfare as well as the ratios such as $\Delta CS_{IP}/\Delta CS_{PF}$ depend on a host of underlying model features, including the price parameter α , the substitutability of products in the demand model (σ), the ability of investors to forecast quality at the time of investment, and the magnitude of the enlargement of the choice set (the share of status quo products available in the higher-cost counterfactual—one-third in the default).

A. *The Price Parameter*

While the price parameter α has no effect on the ΔCS ratio, it has a direct effect on the absolute measure ΔCS_{IP} . Our baseline α is a bound derived from assuming revenue-maximizing song pricing. If we instead assumed that prices were set such that the elasticity of demand were one-half rather than one, then α would be half as large, and ΔCS_{IP} would double. By contrast, if prices were set such that the elasticity were two, then ΔCS_{IP} would be half as large.²⁷

²⁷ In Aguiar and Waldfogel (2016b) we also consider how our estimates vary with the share of products included in the counterfactual choice set, concluding that our random long tail in production is substantially larger than the conventional long tail for a wide range of choice set enlargements. We also explore the sensitivity of the ΔCS ratio to alternative demand models.

B. Substitution Parameter σ

Each value of σ gives us a new vector of product qualities δ , which we term $\delta(\sigma)$. Each new δ vector, in turn, can be used to construct forecasts of expected quality. We can use these to create estimates of ΔCS_{IP} and the ΔCS ratio to see the sensitivity of these measures to σ . As illustrated in Section V.A.3, the absolute change in ΔCS_{IP} depends on the substitution parameter σ . Using our baseline σ of 0.751, ΔCS_{IP} is \$10 million. By contrast, if σ were at the ends of our estimated ranges (0.5 or 0.9), then ΔCS_{IP} would be about \$13 million or \$5 million, respectively.

It is not clear a priori how different levels of substitution affect the ΔCS ratio, so we undertake simulations for different values of σ . For each possible σ , we calculate a δ vector, perform our two-step forecast, and then calculate the ΔCS ratio. Figure 2 depicts the relationship between σ and the ΔCS ratio. Our estimate of the ratio is nearly invariant to our choice of σ . If $\sigma = 0$, then this becomes the plain logit model, and the ΔCS ratio is 19.78; if $\sigma = 0.9$, the ratio is 19.88. Using the potentially overfitted but conservative 2011 forecasting model, this ratio is nearly constant at 13. Because σ is the only estimated parameter determining δ , figure 2 also contains implicit estimates of the standard error of our ΔCS ratio estimate. That the ΔCS ratio is nearly invariant in σ means that if we take bootstrap draws from the estimated σ distribution, the resulting values of the ΔCS ratio would be tightly distributed. We conclude that our estimates of the




FIG. 2.— ΔCS ratio and substitution parameter σ

ΔCS ratio are not sensitive to the choice of logit versus nested logit, nor are they sensitive to the degree of substitutability among songs. Beyond this, our estimates of the ΔCS ratio are precise.

C. Investors' Forecasting Ability

One of the key features of the model is the extent to which investors can forecast quality at the time of investment. The better their ability to forecast, the smaller are both ΔCS_{IP} and the ΔCS ratio. Hence, we would like to investigate the sensitivity of our welfare estimates to different abilities to forecast, which we measure through R^2 .

Ideally, we would like to see beyond the veil of our ignorance to understand how our forecasting ability improves as we add more variables. Of course, we have already included all of the variables available to us in our forecast. To see how our estimate would change if we had better ability to forecast, we can create a new explanatory variable that is the true value of δ plus a scaled random error. That is, define $B_j = \delta_j + sv_j$, where s is a scaling variable that we control and v is a standard normal error.

Then our forecasting model regresses δ on x_{jc} as in column 4 of table 3 in Section V.B above, along with B . We begin with a large value of s , so that we are adding an irrelevant variable, whose coefficient will be small.²⁸ As s shrinks, B acquires a significant coefficient and our ability to predict quality improves. Each value of s is thus associated with a regression R^2 and a prediction R^2 . Figure 3 depicts the relationships between ΔCS_{IP} and the ΔCS ratio and the prediction R^2 .

When $s = 1,000$, the regression R^2 rises from its baseline of .403 to .495, and the associated prediction R^2 (for vintage 2011 alone) rises from its baseline of .323 to .434. The ΔCS ratio would fall from its baseline value of nearly 20 to about 7.6, and the absolute change ΔCS_{IP} would fall from its baseline of \$10.09 million to about \$3.9 million. When $s = 250$, the regression R^2 rises to .846 and the associated prediction R^2 is .843.

A conservative approach to measuring the welfare gain from digitization would employ the richest and most accurate prediction model available. One model that errs on the side of conservatism is the model estimated on 2011 data so that the regression residuals are direct "forecasts" of quality (as opposed to using the forecasts derived from the 2010 regression). That regression had an R^2 of .411. The resulting estimates of ΔCS_{IP} and the ΔCS ratio are \$6.63 million and 13.03, respectively.

²⁸ We use the following values of s to perform our exercise: 10,000,000, 5,000,000, 1,000,000, 500,000, 100,000, 10,000, 1,000, 900, 800, 700, 600, 500, 450, 400, 350, 300, 250, 200, 150, 100, 50, 10, 5, 1, and 0. A value of $s = 10,000,000$ gives rise to our baseline prediction R^2 of .325.

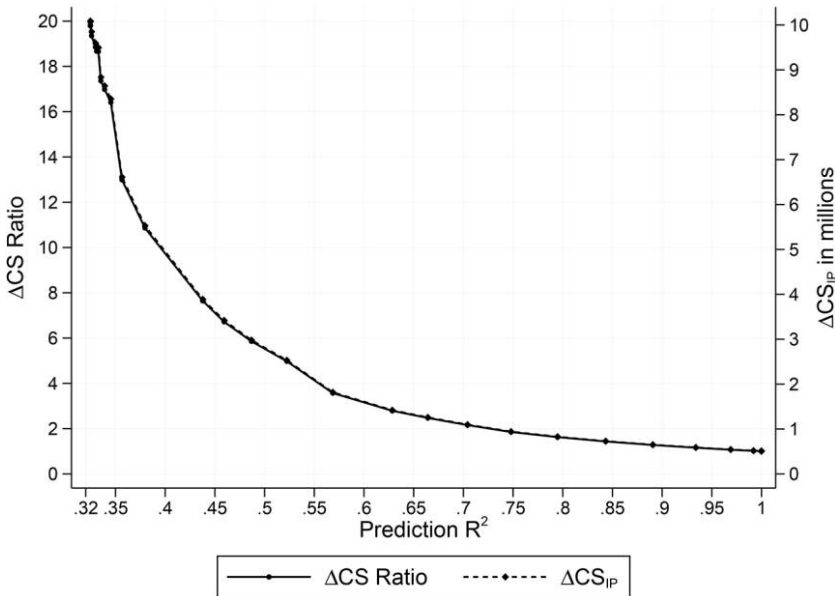


FIG. 3.— ΔCS ratio, ΔCS_{IP} , and prediction R^2

VIII. Conclusion

Evaluating the benefit of new products is a central task for economics. Our study has three conclusions. First, unpredictability can have a large effect on the impact of new product entry on welfare. We explore the welfare benefit arising from the new products prompted by reduced entry costs in a context in which quality is unpredictable. This unpredictability has a large effect on the benefit of new products. Given that unpredictability appears to be a common feature of new products, this idea may have wider applicability.

Second, applying this perspective to the impact of digitization on the recorded music industry yields some novel insights about the benefit of the internet. Observers have understood the benefit of the internet to operate through a shelf space mechanism that we have termed the long tail in consumption. As important as this mechanism is, we propose that the long tail in production that we explore is quantitatively more important. Reductions in entry costs allow producers to “take more draws,” and given the unpredictability of quality at the time of investment, taking more draws can generate more “winners.” Our illustrative estimates for music show that the production mechanism could generate almost 20 times as much benefit as the consumption mechanism for an equal-sized increase in the number of products. This is invariant to the demand estimates and in-

stead depends on the predictability of product quality. It is hard to know the exact predictability of quality, but given the evidence—here and elsewhere—on the unpredictability of the commercial appeal of cultural products, it seems safe to say that the random long tail is likely to be substantially larger than the conventional one. Unpredictability is a generic feature of creative products such as books and movies, suggesting that the growth of new products in those categories may be producing large welfare benefits (Waldfogel and Reimers 2015; Waldfogel 2016).

Finally, the results of this study provide evidence of an explicit mechanism by which the growth in new music products since Napster has raised the realized quality of music, as Waldfogel (2012) and Aguiar and Waldfogel (2016a) have argued, despite the collapse of recorded music revenue.

References

- Aguiar, L., and J. Waldfogel. 2016a. "Even the Losers Get Lucky Sometimes: New Products and the Evolution of Music Quality since Napster." *Information Econ. and Policy* 34:1–15.
- . 2016b. "Quality Predictability and the Welfare Benefits from New Products: Evidence from the Digitization of Recorded Music." Working Paper no. 22675 (September), NBER, Cambridge, MA.
- Anderson, C. 2006. *The Long Tail: Why the Future of Business Is Selling Less of More*. New York: Hyperion.
- Berry, S. T. 1994. "Estimating Discrete-Choice Models of Product Differentiation." *RAND J. Econ.* 25 (Summer): 242–62.
- Berry, S. T., and P. Haile. 2015. "Identification in Differentiated Products Markets." Working Paper no. 21500 (August), NBER, Cambridge, MA.
- Berry, S. T., J. Levinsohn, and A. Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica* 63 (July): 841–90.
- Berry, S. T., and J. Waldfogel. 1999. "Free Entry and Social Inefficiency in Radio Broadcasting." *RAND J. Econ.* 30 (Autumn): 397–420.
- . 2016. "Empirical Modeling for Economics of the Media: Consumer and Advertiser Demand, Firm Supply and Firm Entry Models for Media Markets." In *Handbook of Media Economics*, vol. 1A, edited by Simon P. Anderson, Joel Waldfogel, and David Strömberg. Amsterdam: Elsevier Sci.
- Björnerstedt, J., and F. Verboven. 2013. "Merger Simulation with Nested Logit Demand—Implementation Using Stata." Konkurrensverket Working Paper Series in Law and Economics, no. 2013:2 (April), Swedish Competition Authority, Stockholm.
- Brynjolfsson, E., Y. J. Hu, and M. D. Smith. 2003. "Consumer Surplus in the Digital Economy: Estimating the Value of Increased Product Variety at Online Booksellers." *Management Sci.* 49 (November): 1580–96.
- Caves, R. 2000. *Creative Industries: Contracts between Art and Commerce*. Cambridge, MA: Harvard Univ. Press.
- Danaher, B., Y. Huang, M. D. Smith, and R. Telang. 2014. "An Empirical Analysis of Digital Music Bundling Strategies." *Management Sci.* 60 (6): 1413–33.
- Ferreira, F., A. Petrin, and J. Waldfogel. 2013. "Trade, Endogenous Quality, and Welfare in Motion Pictures." Working paper (March), Wharton School, Univ. Pennsylvania.

- Gentzkow, M., and J. M. Shapiro. 2010. "What Drives Media Slant? Evidence from U.S. Daily Newspapers." *Econometrica* 78 (1): 35–71.
- Goldman, W. 1984. *Adventures in the Screen Trade*. New York: Grand Central.
- Gourville, J. 2005. "The Curse of Innovation: A Theory of Why Innovative New Products Fail in the Marketplace." Working Paper no. 06-014, Harvard Bus. School.
- Handke, C. 2012. "Digital Copying and the Supply of Sound Recordings." *Information Econ. and Policy* 24 (1): 15–29.
- Hausman, J. A., and G. K. Leonard. 2002. "The Competitive Effects of a New Product Introduction: A Case Study." *J. Indus. Econ.* 50 (3): 237–63.
- Nevo, A. 2000. "A Practitioner's Guide to Estimation of Random-Coefficients Logit Models of Demand." *J. Econ. and Management Strategy* 9 (4): 513–48.
- Oberholzer-Gee, F., and K. Strumpf. 2010. "File Sharing and Copyright." In *Innovation Policy and the Economy*, vol. 10, edited by Josh Lerner and Scott Stern, 19–55. Chicago: Univ. Chicago Press (for NBER).
- Petrin, A. 2002. "Quantifying the Benefits of New Products: The Case of the Minivan." *J.P.E.* 110 (4): 705–29.
- Quan, T., and K. Williams. 2017. "Product Variety, Across-Market Demand Heterogeneity, and the Value of Online Retail." Discussion Paper no. 2054, Cowles Found. <http://dx.doi.org/10.2139/ssrn.2871513>.
- Shiller, B., and J. Waldfogel. 2011. "Music for a Song: An Empirical Look at Uniform Pricing and Its Alternatives." *J. Indus. Econ.* 59 (4): 630–60.
- Sinai, T., and J. Waldfogel. 2004. "Geography and the Internet: Is the Internet a Substitute or a Complement for Cities?" *J. Urban Econ.* 56 (July): 1–24.
- Terviö, M. 2009. "Superstars and Mediocrities: Market Failure in the Discovery of Talent." *Rev. Econ. Studies* 76 (2): 829–50.
- Waldfogel, J. 2012. "Copyright Protection, Technological Change, and the Quality of New Products: Evidence from Recorded Music since Napster." *J. Law and Econ.* 55 (4): 715–40.
- . 2015. "Digitization and the Quality of New Media Products: The Case of Music." In *Economic Analysis of the Digital Economy*, edited by Avi Goldfarb, Shane Greenstein, and Catherine Tucker. Chicago: Univ. Chicago Press (for NBER).
- . 2016. "Cinematic Explosion: New Products, Unpredictability, and Realized Quality in the Digital Era." *J. Indus. Econ.* 64 (4): 755–72.
- Waldfogel, J., and I. Reimers. 2015. "Storming the Gatekeepers: Digital Disintermediation in the Market for Books." *Information Econ. and Policy* 31:47–58.