Visibility Incentives for Sellers in a Digital Goods Marketplace

Justin Huang*

September 12, 2017

Abstract

There has been rapid growth in digital goods platform-marketplaces such as the Apple App Store or Google Play store, partially driven by low barriers to entry for new and entrepreneuring developers looking to sell their products. These small and independent sellers rely on marketplace mechanisms for product discovery in the absence of the means or ability to advertise. App store markets frequently link visibility and discovery to historical sales – top selling products are featured more prominently on frontpage and search results. Similar mechanisms are present in many online store contexts. Yet, their impact on marketplace outcomes has largely been understudied in the literature. This research aims to bridge this gap using data from a large app store-like marketplace to examine the implications of this common marketplace visibility policy. I present a model in which the visibility policy is internalized by sellers in their pricing decisions and incorporate rich institutional details from the marketplace to aid in estimation. Using this model, I show that marketplace policy leads sellers to compete for visibility by adopting strategies of introductory pricing. Introductory pricing reduces aggregate seller revenue and thereby the marketplace operator's revenue share, as prices are below optimal on average. Furthermore, the highest valuation early buyers are offered the lowest prices. This issue can be thought of as a form of channel conflict from the perspective of the marketplace operator, whereby the operator in the role of the wholesaler would like to enforce a higher price level but is thwarted by seller (retailer) competition for visibility. I evaluate a mitigation strategy the marketplace operator could pursue - ranking products by revenue rather than sales quantity - through counterfactual simulation.

*Justin is a PhD Candidate at Stanford Graduate School of Business. He can be reached at jthuang@stanford.edu.
This is an incomplete working paper that reports only preliminary findings. Please do not cite or circulate.

1 Introduction

The platform-marketplace model has experienced rapid growth as a paradigm for the production and sale of digital goods. Two prominent examples in the Apple App Store and Google Play stores grew from \$10 billion and \$1 billion revenue in 2013 to \$29 billion and \$6 billion in 2016.[1] In this paradigm, the platform provides developer tools and frameworks while outsourcing the actual production of digital goods to independent developers who create software that is sold on an online marketplace. Sellers in these contexts are numerous but often limited in scale and means. In the absence of the ability to advertise, sellers must rely on the marketplace in order to reach an audience for their goods. Thus, marketplace policies that allocate buyer attention have enormous impact on seller profitability.[2] A common policy in these marketplaces is to give higher visibility to top-selling products. I refer to a policy of allocating visibility on the basis of historical sales as a "sales-ranked" visibility policy. This is functionally implemented through the use of recommendation systems, frontpages, and positioning in search rankings. An example is the "Top Charts" feature on the Apple App Store, which displays a ranked list of apps based on recent sales. Sales-ranked visibility policies allow commonly purchased products to reach a larger audience and encourage sellers to create the kinds of products buyers desire. There are other consequences, however, to this ranking policy. Sellers in online marketplaces typically have a high degree of flexibility over pricing, and the linkage of sales and future visibility encourages sellers to compete for visibility through discounting strategies such as introductory pricing and price pulsing. To what extent does this affect outcomes on the marketplace? I evaluate a sales-ranked visibility policy using data from an app store-like marketplace and a model that rationalizes the price discounting strategies observed.[3]

This paper uses a unique dataset obtained from partnership with an app store-like platform-marketplace. The empirical context is exceptionally well-suited to the study of new product pricing dynamics and platform visibility policies. New product pricing strategy is usually difficult to study due to infrequency of new product releases, pricing arising from multiple agents' incentives

[1] http://bgr.com/2013/11/19/google-play-annual-revenue/,
https://www.statista.com/statistics/296226/annual-apple-app-store-revenue/,
http://www.businessinsider.com/stats-on-googles-revenues-from-youtube-and-google-play-2015-7

[2] Buyer attention and visibility can be taken as synonymous for the purposes of this paper.

[3] For clarity, this paper will refer to the technological platform, the set of software and/or hardware tools which form a base upon which other applications and technologies are developed, as the *platform*. The *marketplace* is the e-commerce store which serves as a centralized clearing house connecting buyers and sellers of technology products related to the platform. Marketplace roles include processing transactions, mediating disputes, tracking reputation, and allocating visibility. In this paper's context, the platform and marketplace operators are the same firm, similar to Apple's development of the iOS operating system and iPhones and their corresponding Apple App Store for sale of iOS applications.

within the supply channel, requirement to infer potential market size, and unknown intertemporal demand linkages. In my context, these concerns are largely resolved. First, I observe a very large number of new product introductions. Second, sellers on the platform have high flexibility in setting their desired prices. Finally, I observe day-level impressions for each product and know the algorithm used by the marketplace for implementing sales-ranked visibility policy. I begin by documenting systematic price patterns in the data. Sellers on the marketplace tend to pursue a strategy of introductory pricing characterized by low yet increasing prices for newly launched products followed by a long-run price decline. A subset of sellers even engages in price pulsing (i.e. alternating between high and low prices). I propose a model that rationalizes these pricing strategies through dynamic incentives arising from the sales-ranked visibility policy on the marketplace. In this model, the marketplace tracks product-specific ranking scores that summarizes a product's historical sales into a single scalar-valued state. This ranking score acts as a shifter for a product's position in search rankings, allocating more buyer impressions to products with higher ranking score. Forward-looking sellers have full knowledge of the ranking score evolution and its effect on future visibility and account for these dynamics in choosing an optimal pricing path. I show that sales-ranked visibility policy induces sellers to compete for visibility through steep price cuts, reducing aggregate seller revenue and thereby the marketplace's revenue cut. I estimate this model and use parameter estimates to evaluate counterfactual mitigation policies the platform could pursue, including allowing sellers to buy advertising and alternate ranking policies.

In this paper, I present a novel mechanism through which sellers introduce new products at a discounted price or indulge in price pulsing. I demonstrate that firms have incentives for such pricing strategies even with homogenous, myopic consumers, specifically in the context of platform marketplaces, due to the linkage between prices and future visibility. A large literature in marketing and economics has studied price dynamics in the context of durable goods theoretically (Coase [1972], Stokey [1979], Moorthy [1988], Narasimhan [1989], Balachander and Srinivasan [1998]), and to a more limited extent, empirically (Nair [2007], Dube et al. [2009]). This literature focuses on price skimming, where firms charge higher prices to early customers with high willingness to pay and low patience, and lower prices to later customers who anticipate lower prices in the future and choose to wait to purchase. By contrast, I study introductory pricing where firms offer a lower price in early periods in order to gain visibility in the future. This mechanism exists even with homogenous, myopic consumers. A large literature on price promotions or sales has focused on these as means of discriminating between heterogeneous consumers (Salop [1979], Varian [1980], Conlisk et al. [1984]). Another literature has focused on introductory price discounts due to the quality signaling effect of prices (Milgrom and Roberts [1987]), the cost signaling effect (Bagwell [1987]), and installed base effects (Farrell and Saloner [1986], Katz and Shapiro [1992]). I demonstrate that pricing dynamics can exist even in the absence of heterogeneity in valuations, search costs, reservation prices, signaling mechanisms, and installed base effects mentioned above. While my discussion is in the context of platform-marketplaces, the ideas are broadly applicable to any environment where current sales influence future demand, for example in the presence of word-of-mouth effects.

This paper joins a broad literature on topics surrounding platform-marketplaces, including competition (Rochet and Tirole [2003]), pricing (Weyl [2010]), reputation systems (Nosko and Tadelis [2015], Cabral and Hortacsu [2010]), and search (Yao and Mela [2011], Fradkin [2014], Tadelis and Zettelmeyer [2015]). The literature on platform pricing has focused on the operator's pricing decision rather than those of sellers on the platform-marketplaces. In economics and marketing, the platform search ranking literature has considered rich models of buyer behavior while generally considering supply invariant to ranking policies. The search ranking literature in information systems and computer science (Chapelle and Zhang [2009] and Joachims [2002]) has a similar buyer-side focus, aiming to optimize search ranking results using a massive set of page features, query history, and click logs. These algorithms typically have an objective of maximizing buyer clickthrough or expected revenue from the search with consideration for quality uncertainty and heterogeneity. To my knowledge, this is the first paper to examine marketplace search ranking design in the presence of dynamic incentives for sellers or content creators. I demonstrate that the commonly-used sales-ranked visibility policies in the industry may lead to poor market outcomes in the presence of endogenous seller reactions. The finding of inefficiency in platform search ranking design bears similarity a few works, including: Arnosti et al. [2014], who illustrate how ranking systems promoting top sellers exacerbate congestion; Bresnahan et al. [2015], who examine a similar app store context where sellers have ranking incentives to purchase downloads from third party sites; and Choi and Mela [2016], who model seller advertising choice in sponsored search rankings with a rich buyer-side search model. Lastly, this paper is related to the small but growing literature on the incentives for content creators in social network and other platform contexts, including Huang and Narayanan [2016], Zhang and Zhu [2011], Sun and Zhu [2013] and Toubia and Stephen [2013].

I show that the seller endogenous reaction is consequential in the design of search rankings through simulations on model estimates. A sales-ranked visibility policy that links historical sales to future visibility creates an incentive for sellers to compete for visibility through price discounting strategies. The discounting incentive is strongest for both new products and high-quality products due to increased likelihood of transitioning to the highest and most profitable ranking states. Discounting behavior reduces aggregate seller revenue compared to a counterfactual scenario where sellers set price without internalizing ranking dynamics, such as if sellers were unaware of the mechanism or the marketplace could mandate prices. The reduction of aggregate seller revenue translates directly into a loss of revenue for the platform-marketplace operator, since the operator takes a fixed percentage cut of all transactions on the marketplace. Reduction of aggregate seller revenue arises from two related but distinct routes: price levels are lower than optimal due to competition for visibility, and incentives for introductory pricing undermine effective intertemporal price discrimination. The platform ultimately accepts lower than optimal prices to shuffle a finite amount of visibility among competing products. This inefficiency can be likened to channel coordination conflict, whereby a wholesaler would like to mandate retail price floors but channel profitability is ultimately undercut by retail competition (Cachon [2003]). Two other counterfactual scenarios are evaluated, including a ranking score based on sales revenue rather than sales quantity and allowing sellers to buy advertising on the marketplace [In Progress].

The rest of the paper proceeds as follows: Section 2 describes the paper's empirical context, an app store-like marketplace for the distribution of platform-specific digital goods. Section 3 examines the data and descriptive patterns which will motivate modeling decisions. Section 4 presents a dynamic structural model of pricing and purchase on the platform. Section 5 details model estimation and results. Section 6 [In Progress] discusses managerial implications and evaluates a number of counterfactuals based on model parameters. Section 7 concludes.

2 Empirical Context

This section introduces and describes the empirical context for the paper, including a motivation for the choice of context, marketplace description, roles of buyers and sellers, and a discussion of search rankings.

The goal of this paper is to examine the effects of a commonly-employed search and visibility policy of platform-marketplaces on seller incentives and market outcomes. As described earlier, there is a linkage between historical sales and position and visibility in search rankings. I will refer to this as a "sales-ranked" visibility policy. I collaborate with a firm that operates a large marketplace for software goods. This marketplace shares relevant institutional features and details with app stores such as Google Play or the Apple App Store, including sales-ranked visibility, a large number of software products, and a high degree of seller pricing flexibility. The empirical context for this study is an app store-like marketplace for digital goods. This marketplace provides a centralized point for the sale and distribution of goods, manages seller reputation, and curates search results. The goods in this context are licenses for software exclusively usable on a technology platform which the company also manages. Once purchased, the license has no expiration and entitles the buyer to unlimited use of the underlying software code. Thus, this is a durable digital good with no repeat purchase and no in-app purchase or other revenue streams for the seller once the sale has been made. The online frontpage of the marketplace features top selling and recommended products and prominently displays a search bar.

The thousands of sellers on this marketplace create and sell middleware which aids other developers. This middleware is named as such because it bridges low-level routines provided by the platform and high-level applications. Middleware consists of software packages that implement simulation and visualization tools to simplify and accelerate software development on the platform. The millions of buyers on the marketplace are small to medium-scale software developers who purchase middleware. The distinction between sellers and buyers on the marketplace is in their audience and the types of software products they produce. Sellers are creators of middleware whose audience is other software developers, while buyers purchase and utilize middleware in order to create finished applications for sale to mass-market consumers outside of the marketplace.

Sellers have the ability to set and change prices for their products as frequently as desired. There are no competing marketplaces for the software sold on this platform, and advertising or sales of these software products outside of the marketplace is nearly non-existent. Further, there is no ability for sellers to buy advertising on the platform itself. Thus, the platform has a high degree of control over the sale and visibility of these digital goods through levers such as site design, search rankings, and frontpage mechanisms.

Search ranking is the primary tool utilized by this marketplace to allocate visibility of products and thereby match buyers to sellers. For a given search query, the rankings of products returned arises from the interaction of a product-specific ranking score and the relevance of the product attributes to the query. Through the collaboration with the firm, I have specific information about the ranking score and how it affects search results. The ranking score is scalar-valued and does not vary by query or keywords of the search. It summarizes the history of sales and other factors relevant to the quality of the product into a single score. This ranking score is product-specific rather than seller-specific, and there are no additional factors that enter search rankings. The evolution of this ranking score is known and deterministic, combining inputs related to buyers' interactions with the products and other factors. While the full details of the evolution of the ranking score cannot be shared, for the purposes of this paper it can be thought of as primarily a function of historical sales quantity - this will be described in greater detail in Section 4. I incorporate these details about the ranking score into the model.

3 Data and Descriptives

This section describes the dataset used in this paper. As previously mentioned, this dataset arises from a research collaboration with a large technology platform-marketplace operator.[4] In this section, I first describe the data, and then explore empirical patterns in pricing and sales to document the linkage between the platform's visibility policy and pricing decisions that are induced by it. These empirical regularities form the basis for the structural empirical model I develop and estimate.

3.1 Data

The dataset for this paper comes directly from the collaborating platform-marketplace operator. This dataset spans a roughly six year period from December 2010 through March 2017. It includes event-level data for all products listed on the marketplace, including a product identifier, listing timestamp, price, category, keyword tags, description, and price change history. On the buyer side, the dataset includes an event-level purchase history encompassing every sale on the platform and timestamped product reviews including star rating and text. In addition, I have day-level counts of unique buyer impressions for each listed product. All of these datasets are collapsed and merged to a product-day level for analysis.

3.2 Descriptives

The descriptives presented in this section document empirical patterns in pricing and sales on the platform. As a preview of this descriptive analysis, I present a few main empirical regularities as follows.

1. The overall price trend involves lowered introductory price, followed by an increase and then steady decrease to below starting levels.

2. A small subset of sellers pursue a strategy of price pulsing between high and low price levels, with the low levels including a $0 price.

3. The sales pattern shows an initial ramp up and then a long run decline similar to the price pattern. Peak sales occur later than peak prices.

[4] In order to protect the privacy of the data source and at their request, figures relating to aggregate user counts, sale quantities, prices and revenues have been normalized.

4. Demand is decreasing in price.

5. Previous period sales increase current period visibility as measured by product impressions.

Figure 1 shows the price path of products listed on the marketplace, charting their mean price levels over time. Prices are normalized by dividing by the mean price of all products on day 1. Figure 1 provides mean price levels and error bars documenting a 95% confidence interval for the mean. Prices are seen to increase by approximately 6% for the first 15 days after product launch and then steadily decline towards a long-run level approximately 5% lower than the starting price level. Regressions shown in table 1 of prices in days 1-15 and 16 and onwards with product level fixed effects demonstrate that this pricing pattern is significant on an individual product level in addition to in aggregate. Sellers of new products seem to be pursuing a strategy of lowered introductory pricing on the marketplace.

Figure 2 shows an example to support the notion that ranking effects are known to sellers and internalized in their pricing decisions. The seller whose product price trajectory is displayed here is an experienced seller on the platform. The seller can be seen to employ a series of $0 prices near product launch and various other periods punctuating a typical price level of $25-$50. The $0 price level would clearly be suboptimal for a profit maximizing seller in the absence of ranking effects, and thus is taken as an indication that the seller intends to attain sales during these low price periods in order to boost rankings.

The sales pattern can be seen in Figure 3. Sales tend to increase over the first 25 days to a maximum approximately 200% of first day sales. From there, sales are seen to decline by roughly a third from peak levels by the end of the 180-day observation period shown. Regressions with product fixed effects in Table 2 confirm these effects on the individual level. Peak sales occur later than peak prices, which is consistent with a price strategy which internalizes ranking effects if individual product pricing is taken to be representative of the aggregate trend. Pricing has an immediate effect on sales through buyer preferences and a delayed effect through increased search ranking. The later sales peak allows sellers to capitalize on the costly strategy of introductory pricing and harvest a significant fraction of sales at peak or near-peak prices.

Regressions of prices on sales quantity are shown in Table 3. Without controlling for product fixed effects, the effect of price on sales is positive, likely due to differences in product quality and the fact that both prices and sales correlate with product age as shown in Figures 1 and 3, rising initially and falling in the long run. Controlling for product fixed effects, the sign of price effect reverses to negative. The addition of a polynomial control for ranking score dramatically increases adjusted R^2 without significantly changing the price effect. That the effect of price on sales is negative and significant lends credibility to a seller pricing strategy that would cut prices in order to increase rankings.

Finally, the regression in Table 4 shows the relationship between previous period sales and current period visibility as measured by product impressions. At the day level, I regress current period product impressions on lagged unnormalized sales quantity and lagged product impressions with product fixed effects. If interpreted causally, the highly significant coefficient on lagged sales suggests that 1 additional sale generates 1.2 additional product impressions in the subsequent period. Daily product impressions are seen to have a strong and significant serial correlation.

3.3 Discussion

I have documented the presence of both introductory pricing and long-run price skimming in the marketplace. These pricing patterns hold not only across products, but within products as well. Introductory pricing suggests that sellers have an incentive to generate early sales, such as in the presence of search ranking dynamics, word of mouth, or installed base effects. I argue that the alternative economic motivations for introductory pricing can be ruled out by the empirical context. Word of mouth surrounding these products is scant, as the middleware products sold are niche and rarely discussed outside the marketplace. On the marketplace itself, communication between buyers is only possible through reviews which are captured in the dataset. Reviews are quite rare - occurring in roughly .1% of purchases with less than 5% of products having at least one review. Further, the inclusion of reviews in modeling shows little effect on purchase when controlling for previous sales quantity and ranking score. Installed base effects are not relevant for the software sold, as there is no network or collective benefit to other buyers using the same middleware product. This leaves ranking dynamics as a plausible motivation for introductory pricing. That said, since these other phenomena lead to similar incentives for lower introductory prices, my model has broader applicability to other empirical contexts.

Taken together, the results of Table 3 and Table 4 demonstrate the viability of a pricing strategy incorporating visibility dynamics. Downward sloping demand implies that a seller could use price discounts to increase sales. These sales are valuable to a forward-looking seller beyond the revenue they generate, as sales boost future product visibility through the mechanism of search rankings and the ranking score. The next section will present a model which captures the seller dynamic incentives induced by the sales-ranked visibility policy and rationalize the patterns of pricing seen in the data.

4 Model

I present a stylized model of the marketplace which captures the strategic incentives of forward looking sellers. The model assumes forward-looking sellers with complete information about demand, the effect of ranking score on visibility, and the ranking score evolution process.[5] Sellers in the model make pricing decisions in each period that internalize the incentive structure induced by sales ranked search results: higher current period sales will improve rankings and thereby product visibility in future periods. The visibility of a product, which is observed through product level daily impression counts in the data, is operationalized by a market size of buyers that consider the product in a given period. The seller price decision is driven fundamentally by the tradeoff between increased prices yielding higher revenue per sale and lowered prices yielding increased sales quantity which could be considered an investment in future visibility through higher rankings. On the demand side, buyers are assigned to products by the marketplace search algorithm based on product attributes and ranking score. Buyers are assumed to be myopic and homogenous, making a one-shot decision to purchase or not purchase based on product attributes, posted price, and product age.[6]

4.1 States and Decisions

Let π denote seller payoff for setting a price P_{jt} at a given state S_{jt} where j indexes products and t indexes time as measured in days from product launch, and let δ be the assumed discount factor. I specifically denote this as seller payoff rather than profit because I allow for unobserved and monetary and non-monetary payoffs to sellers outside of the platform. Each seller looks to maximize the expected sum of long-run discounted payoffs, $E[\sum_{t=1}^{\infty} \delta^t \cdot \pi(P_{jt}, S_{j,t})]$ which comprises of profits derived from sale of products on the marketplace, transaction costs associated with changing price, and outside benefit derived from selling goods on the market. Sellers face demand for products and transaction costs that are dependent on the states S_{jt}, which is comprised of the ranking score r_{jt}, previous period price $P_{j,t-1}$, and age of the product in days, denoted t. Ranking score r_{jt} affects demand through current period market size M_{jt}. Previous period price $P_{j,t-1}$ is a relevant state because there are seller transaction costs associated with changing prices. Product age t is included as a state variable because of its impact on demand. Sellers choose the optimal price path $\{P_{jt}\}_{t=1}^{\infty}$ of each product under their control as monopolists with full knowledge of their product's current state S_{jt}, the distribution of states of other products on the marketplace that is assumed to be

[5]The assumption on complete information about demand will be relaxed in future work

[6]This will be relaxed in future work

fixed, the marketplace algorithm for allocation of visibility operationalized as market size M_{jt}, and the demand function $q(P_{jt}, S_{jt})$ induced by buyer utility. Buyers on the platform are myopic and homogenous, making a one-shot decision to purchase or not purchase the product they are shown by the marketplace.

4.2 Model Timing

The timing of the model is structured as follows:

1. At the start of each period t, each package j is endowed with a market of homogenous buyers $i \in I_{jt}$ such that the count of buyers $|I_{jt}| \equiv M_{jt} = \psi(x_j, \tilde{r}_{jt})$, where x_j are product characteristics, $\tilde{r}_{jt}$ is the ordinal of the ranking score r_{jt}, and ψ approximates the search process on the marketplace.

2. Sellers observe the market size M_{jt} and set a price P_{jt} for each owned package j based on forward looking payoff maximization incorporating knowledge of demand, current market size and ranking score, previous price level and associated price change transaction cost, and the state evolution process. The seller is assumed to have complete knowledge of the states and evolution processes.

3. Each buyer i allocated to product j observes the chosen price P_{jt} and chooses between non-purchase and purchase $y_{ijt} \in \{0, 1\}$ according to the specified buyer utility function $U(P_{jt}, x_j, S_{jt})$.[7] Buyers are assumed to maximize this utility in their decision.

4. The marketplace observes buyer decisions and updates the ranking score state $r_{j,t+1}$ according to the known rank score evolution function.

4.3 Search Ranking Allocation

The search ranking allocation is a policy set by the platform to allocate search rankings and thereby visibility from the large number of buyers entering the marketplace and placing search queries. The number of impressions for each product j on day t is observed in the data and operationalized in the model as a product-day market size M_{jt}. This market size varies significantly across products

[7] Buyers observe the state space partially. While they observe and are assumed to care about the age of products t, they do not observe past prices $P_{j,t-1}$ or rank r_{jt}. This is reasonable in this context because monitoring past prices would require searching on a regular basis, and observing ranks would additionally require observing the position of a particular product across multiple queries.

and within products over the course of their lifetime. The market size is modeled as follows:

$$M_{jt} = \psi(x_j, \tilde{r}_{jt}) \tag{1}$$

where ψ is a function of product attributes x_j and the ordinal ranking of the product's ranking score $\tilde{r}_{jt}$. This allocation of buyers can be thought of as an approximation of a search rankings process in which: a product's position in search rankings for a given query is a combination of the ranking score score and the distance between product attributes and the specified query; buyers only consider position when deciding which product to view; and the space of potential queries is large and constant over time.[8] The market size allocation is flexibly estimated with a log-linear regression including polynomial terms for ordinal ranking score and a rich set of product observables capturing category, keywords, and description text.

The ranking score is product-time varying scalar value which can be calculated from the sales record and other factors. Its evolution algorithm has been provided by the platform and can be summarized as follows for the purposes of the paper:

$$r_{jt} = \omega(r_{j,t-1}, q_{j,t-1}) \tag{2}$$

Ranking score of product j in time t is a function ω of previous period ranking $r_{j,t-1}$ and previous period sales $q_{j,t-1}$. This ranking score evolution process is an algorithm provided by the platform-marketplace. ω is increasing in previous period sales quantity, and the ranking score is further known to decrease in periods with 0 sales, or that $\omega(r_{j,t-1}, 0) < r_{j,t-1}$.

4.4 State Transitions

The first state, ranking score r_{jt}, evolves based on the previously described function ω. The ranking score is endogenous to the firm's previous pricing decisions through demand. Assuming a downward sloping demand curve, the ranking score should be decreasing with respect to historical prices. The second state, $P_{j,t-1}$, is equal to the chosen price in the previous period. This state is relevant in the model due to the transaction costs associated with changing prices specified in seller payoffs. While the marketplace itself does not charge sellers to change pricing, the transaction cost captures the effort necessary to access the seller page and make the requisite change. Lastly, the product's age t as measured from publication is a relevant state due to its inclusion in the buyer demand

[8] In my empirical analysis, I restrict my analysis to a 180-day lifespan for each product. Within this relatively short period of time, this assumption is reasonable.

function. t denotes days since publication and increments by 1 in each period.

4.5 Buyer Demand

Buyer demand is specified with a binary logit choice model. Homogenous buyers, indexed $i \in \{1, .., M_{jt}\}$ where M_{jt} arises from the search ranking allocation defined above, make a static choice between purchase and non-purchase. Buyer i's utility of purchasing product j in time t is given by:

$$U_{ijt} = \beta \cdot P_{jt} + x'_{jt}\gamma + \zeta_{jt} + \xi_{ijt} \tag{3}$$

where β captures marginal utility of wealth, γ captures consumer tastes for product attributes and product age, $x_{jt} \equiv [x_j; t]$ combines the time invariant product attributes x_j and product age t, ζ_{jt} is a product-time taste shock common to all buyers and serves as an econometric error term, and ξ_{ijt} is an individual level product-time taste shock rationalizing choice outcomes. The utility of non-purchase is normalized to mean 0 and given by:

$$U_{i0t} = 0 + \xi_{i0t} \tag{4}$$

Purchase follows the typical logit decision rule and purchase probabilities:

$$y_{ijt} = \begin{cases} 1 & U_{ijt} > U_{i0t} \\ 0 & otherwise \end{cases} \tag{5}$$

$$\mathcal{P}(y_{ijt} = 1) = \frac{log(\beta \cdot P_{jt} + x'_{jt}\gamma + \zeta_{jt})}{1 + log(\beta \cdot P_{jt} + x'_{jt}\gamma + \zeta_{jt})} \tag{6}$$

Point estimates of sales quantity can be obtained by inverting this demand equation: $q^*_{jt}(P_{jt}|S_{jt}, x_j, \beta, \gamma) = \frac{log(\beta \cdot P_{jt} + x'_{jt}\gamma)}{1+log(\beta \cdot P_{jt} + x'_{jt}\gamma)} \cdot M_{jt}$. For the purposes of modeling, demand $q_{jt}((P_{jt}|S_{jt}, x_j, \beta, \gamma)$ is simulated from draws rather than using the point estimate since the purchases are discrete and ranking effects introduce significant nonlinear long-run seller returns to a marginal sale.

4.6 Seller Decision

The seller's decision arises from maximizing the sum of discounted payoffs. Each product is assumed to be a separate market, eliminating direct price competition effects and dependencies in the state space S_{jt} on other product decisions. This assumption, associated limitations, and effect on counterfactuals will be revisited in later discussion. The seller's choice of optimal price path

$\{P_{jt}\}_{t=1}^{\infty}$ is then a separable decision for each owned product j and the seller choose prices as a monopolist. Seller current period payoff is given as follows:

$$\pi(P_{jt}|r_{jt}, P_{j,t-1}, t; x_j) = \alpha_j(P_{jt}) - \lambda_j \cdot I[P_{jt} \neq P_{j,t-1}] + E[q(P_{jt}; r_{jt}, x_j, t)] \cdot P_{jt} \tag{7}$$

where P_{jt} is the chosen price, $E[q(P_{j,t}; r_{j,t}, x_j, t)]$ is the expected demand at the chosen price arising from the demand model, $\lambda_j > 0$ is a heterogeneous transaction cost associated with changing prices, and α_j captures heterogeneous non-monetary incentives to distribute content on the marketplace (see: Toubia and Stephen [2013]). The choice of pricing path maximizes the objective function:

$$\max_{\{P_{jt}\}_{t=1}^{\infty}} \sum_{t=1}^{\infty} \delta^t \cdot \pi(P_{jt}) \tag{8}$$

where δ is an assumed discount factor. For modeling purposes, the decision space is split into a finite set of price levels $P_{jt} \in \{P^1, P^2, ..., P^K\}$ chosen by quantiles of the empirical distribution of price choices on the marketplace.[9]

The seller's solution to this optimization problem balances the benefits to discounting or even \$0 price levels yielding additional sales and thereby increased market size in subsequent periods with higher revenues per sale and consideration for frequently incurring price change transaction costs. This gives rise to the following Bellman equation, where V^* is the unique state value function solution arising from an optimal forward decision path:

$$V^*(r_{jt}, P_{jt-1}, t; x_j) = \max_{P^*} \pi(P^*|r_{j,t}, P_{j,t-1}, t; x_j) + \delta \cdot E[V^*(r_{j,t+1}, P^*, t+1; x_j|P^*)] \tag{9}$$

and the following choice-specific value function:

$$V(P_{jt}|r_{jt}, P_{j,t-1}, t; x_j) = \pi(P_{jt}|r_{jt}, P_{j,t-1}, t; x_j) + \delta \cdot E[V(r_{j,t+1}, P_{jt}, t; x_j|P_{jt})] + \epsilon_{kjt} \tag{10}$$

where ϵ_{kjt} is a Type 1 Extreme Value shock observed by the seller but not the econometrician and shifts the payoff of a given price choice in a given period. The shocks are assumed i.i.d. across choices, products, and time. The expectation $E[V(r_{j,t+1}, P_{jt}, t; x_j|P_{jt})]$ is taken with respect to the set of potential future ranking scores $\{r_{j,t+1}\}$ conditional on the price choice P_{jt} since $r_{j,t+1}$

[9] This discretization allows calculation of choice specific value functions via value function iteration on a grid over the state space.

is a function of demand q_{jt} which is stochastic and integer valued. Given that shocks are Type 1 Extreme Value and writing the state as $S_{jt} \equiv \{r_{jt}, P_{j,t-1}, t\}$, the state value function can be written as a function of the choice specific values (Rust [1987]):

$$\begin{aligned} V^*(S_{jt}; x_j) &= \max_{P^*} \pi(P^*|S_{jt}; x_j) + \\ \delta \cdot \int log(\sum_{k=1}^{K} &exp[V(P^k|S_{j,t+1}; x_j)])dF(S_{j,t+1}, \epsilon_{k,j,t+1}|P^*; S_{jt}; x_j) \end{aligned} \tag{11}$$

Both the state value function and choice specific value function can be calculated using value function iteration. The probability of a given price choice is given by:

$$\mathcal{P}(P_{jt} = P^k|S_{jt}; x_j) = \frac{exp(V(P^k|S_{jt}; x_j)}{\sum_{m=1}^{K} exp[V(P^m|S_{jt}; x_j)]} \tag{12}$$

4.7 Discussion

The above model makes a number of assumptions which I will discuss here. These assumptions are largely driven by the empirical context and stylized facts in the data. Notably, on the supply side, sellers maximize a utility function rather than undergoing pure profit maximization. The inclusion of a price-switching transaction cost is driven by an observation on the widely varying likelihood of switching prices between sellers. A large fraction of sellers are observed to never change the price of the product after it has been posted on the market, yet another subset of sellers alter prices for their products on a weekly or even daily basis. A seller-level heterogeneous transaction cost to changing price helps rationalize these behaviors and can be thought of as the effort associated with logging into the platform, assessing a new price level, and changing price accordingly. Sellers that are highly engaged with the marketplace would face a low transaction cost, while sellers that rarely visit the website and must log in specifically to change price would face a high transaction cost. Another aspect of the supply model is the inclusion of a non-monetary utility component $\alpha_j(P_{jt})$ which is a function of price and operationalized as a price-specific intercept term in estimation. This parameter represents outside incentives to contributing on the platform along the lines of Toubia and Stephen [2013], such as gaining reputation as a developer or intrinsic utility from sharing software tools. Outside utility helps rationalize pricing decisions that are inconsistent with forward-looking profit maximization, such as a fraction of sellers which price at and subsequently never deviate from a $0 level. Since ranking effects are product-specific and not seller-specific, this

\$0 strategy has no spillover effects to other products and is never an optimal pricing path, yet it occurs in nearly 8% of products in the dataset.

Sellers are assumed to have full knowledge of the ranking score evolution process. This choice is motivated by discussions with the platform-marketplace operator and app developers. Among developers in this setting it is considered common knowledge that app store marketplaces reward top-downloaded products with increased visibility. Advice websites prominently advise developers that new apps will suffer in ranking from slow download days[10]. Side markets have even popped up allowing the purchase of (sometimes illegitimate) downloads, as studied by Bresnahan et al. [2015]. Sellers are also assumed to know their ranking score, which they might manually approximate by observing where their products rank in search results containing relevant queries. Finally, I assume the platform-marketplace operator can allocate buyer impressions through search rankings independent of buyer-product match. This assumption follows from the relative the strength of reported position effects in the literature, including Narayanan and Kalyanam [2015] and Goldman and Rao [2014]. Data from other environments suggest less than 5% of searchers will visit the second page of a search result.[11] These results imply position effects dominate the match between buyer preferences and product characteristics, and thus the estimated function ψ which maps product attributes and ranking score to impressions can be used to extrapolate impressions under new ranking score values. Collectively, these assumptions put manageable structure on the search rankings allocation process and simplify the state space for feasible estimation.

The use of a binary logit demand model is motivated by the large scale of the platform-marketplace and data considerations. The goal in this paper's demand modeling is to recover own-price elasticities and demand estimates which will enter seller payoff and inform the choice of price. Each product is assumed a separate monopoly market in the model, as the middleware sold on the platform tends to be specific in purpose and exhibits low substitutability even within category or among products with similar keywords. Further, the number of products on the marketplace numbers is in the thousands, making it computationally infeasible to estimate a differentiated goods choice model such as Berry et al. [1995] in which consumers consider the full- or even a sub-set of alternatives. Individual-level search queries, click data, and search results are unavailable, precluding additional structure on the search model similar to that used in Chen and Yao [2016]. Future iterations of this paper may aim to incorporate buyer persistence and finite type heterogeneity, allowing for a more faithful model of intertemporal price discrimination and skimming behavior on the part of the seller. This would require additional assumptions on the type composition of the

[10] https://www.apptentive.com/blog/2016/05/17/30-steps-to-mobile-app-launch/

[11] https://searchenginewatch.com/sew/study/2276184/no-1-position-in-google-gets-33-of-search-traffic-study

observed market sizes and an associated generating process. For now, changing buyer tastes over time are captured through a coefficient of disutility of product age in the demand model, reflecting that older products may be valued less or considered outdated due to advancements in technology. This coefficient can still generate price skimming behavior, but it imposes a functional form and makes buyer taste composition exogenous to sellers' historical pricing decisions.

5 Estimation and Results

Model estimation proceeds in two steps. In the first step, I estimate the demand parameters based on calculated logit shares and product level attributes using the sales record and impressions data. The second step uses the demand parameters to estimate supply side parameters through simulated maximum likelihood estimation wrapped around nested fixed point calculation of the choice specific value functions.

5.1 Demand Estimation

Demand can be rewritten as a logit shares equation which arises from the buyer utility model:

$$log(s_{j1t}) - log(s_{j0t}) = \beta \cdot P_{jt} + x'_{jt}\gamma + \zeta_{jt} \tag{13}$$

where $s_{j1t} = \frac{\sum_{i \in I_{jt}} y_{ijt}}{M_{jt}}$and $s_{j0t} = 1 - s_{j1t}$. Recall that M_{jt} is observed in the data as product-day level impressions. This logit shares equation is estimated using OLS with a subset of parameter estimates reported in Table 5. Demand estimates imply a typical price elasticity of -0.029, suggesting that sellers generally price in the inelastic region and below the optimum which would be chosen by a monopolist in the absence of dynamics. That sellers price in the inelastic region is consistent with the incentives created by sales-ranked mechanisms, as the future market size expansion effect rewards lower pricing that could increase sales today. The age coefficient is negative, suggesting buyers have decreasing valuation for products over time. This negative age coefficient also rationalizes the observed pattern of long-run price decline after the introductory period.

In demand model estimation, there is usually concern for potential endogeneity of prices. If demand shocks ζ_{jt} are observed by the seller, the seller may set prices in response ζ_{jt}, thereby creating an endogeneity bias in the price coefficient. However, I believe endogeneity bias in the above demand estimation to be minor due to the presence of several mitigating factors in this context. First, obvious drivers for large and/or serially-correlated product-time shocks on the platform are absent. The middleware products sold on the platform are not mass-consumer goods and thus do not exhibit seasonal demand or holiday effects. Second, price variation within product is small relative to price variation across products, so most of the price effect is estimated from persistent differences in log-shares across products rather than potentially endogenous variation within products. Finally, price changes by sellers are infrequent relative to the day-level specification of the demand model, reducing potential correlation between price and ζ_{jt}. For these reasons, I proceed with demand estimation assuming price exogeneity.

5.2 Supply Estimation

The supply side of the model is estimated via iterated simulated maximum likelihood using choice specific value functions calculated by nested fixed point. Choice specific value function calculation requires the distribution of potential demand outcomes $\{q_{jt}\}$ given the price choice P_{jt}, state variables S_{jt}, and product attributes x_j. The demand outcomes are drawn from a simulation using the demand parameters and enter both the seller current period payoff $\pi(P_{jt}|r_{jt}, P_{j,t-1}, t; x_j)$ and continuation value $\delta \cdot E[V(r_{j,t+1}, P_{jt}, t; x_j|P_{jt})]$. I discretize the state space into a grid containing 21 price choice points, 13 rank score points, and 10 age points with linear interpolation to calculate value functions between state points. Outside payoff $\alpha_j(P_{jt})$ is similarly discretized into 20 parameters $\alpha_1, ..., \alpha_{20}$ corresponding to the 21 price choice points and α_0 normalized to zero corresponding to the lowest price level $P^0 = \$0$. I fix δ due to the typical difficulty in estimating discount factors in dynamic models (Song and Chintagunta [2003]) Nested fixed point calculation proceeds via value function iteration until convergence.

I implement $R = 2$ type discrete heterogeneity in seller outside payoffs and price switching transaction costs. Heterogeneous outside payoffs rationalize persistent differences between sellers' chosen price levels and profit maximizing price levels implied by the model. Heterogeneous transaction costs rationalize the wide differences in frequencies of price switching observed in the data. Specifically, sellers incur the following penalties to flow payoff for changing price in a given period:

$$\begin{aligned} &-\lambda_1 \cdot I[P_{jt} \neq P_{j,t-1}], \quad \text{if } r = 1 \\ &-\lambda_2 \cdot I[P_{jt} \neq P_{j,t-1}], \quad \text{if } r = 2 \end{aligned} \tag{14}$$

The proportion of sellers of type $r = 1$ is parameterized as $\rho \in (0, 1)$ and sellers of type $r = 2$ as $1 - \rho$. I collect all the supply side parameters $\{\lambda_1, \lambda_2, \rho, \alpha_{1,1}, ..., \alpha_{2,20}\}$ into a parameter vector

Θ.

Given the calculation of choice specific value functions, I proceed to maximum likelihood estimation. The likelihood of a given price choice P_{jt}^k at by seller j at time t given current parameter guess Θ can be written:

$$\mathcal{L}_{jt}(P_{jt}^k|\Theta; S_{jt}; x_j) = \frac{exp(V(P_{jt}^k|\Theta; S_{jt}; x_j)}{\sum_{m=1}^K exp[V(P^m|\Theta; S_{jt}; x_j)]} \tag{15}$$

Then the likelihood of a seller's full pricing history is:

$$\mathcal{L}_j(\Theta) = \prod_{t=1}^{T_j} \mathcal{L}_{jt}(P_{jt}^k|\Theta; S_{jt}; x_j) \tag{16}$$

Accounting for two-type heterogeneity and denoting $\Theta_r \equiv \{\lambda_r, \alpha_1, ..., \alpha_{10}\}$, a seller's full likelihood is then:

$$\mathcal{L}_j(\Theta) = \rho \left(\prod_{t=1}^{T_j} \mathcal{L}_{jt}(P_{jt}^k|\Theta_1; S_{jt}) \right) + (1-\rho) \left(\prod_{t=1}^{T_j} \mathcal{L}_{jt}(P_{jt}^k|\Theta_2; S_{jt}) \right) \tag{17}$$

And the total log likelihood to be maximized:

$$\mathcal{LL}(\Theta) = \sum_{j=1}^{J} log \left[\rho \left(\prod_{t=1}^{T_j} \mathcal{L}_{jt}(P_{jt}^k|\Theta_1; S_{jt}) \right) + (1-\rho) \left(\prod_{t=1}^{T_j} \mathcal{L}_{jt}(P_{jt}^k|\Theta_2; S_{jt}) \right) \right] \tag{18}$$

This log likelihood function is maximized to recover parameter estimates using a general optimization routine of the stats package in R.

5.3 Results

Demand estimation yields the parameters shown in Table 5. The estimated price parameter of -0.00132 implies a price elasticity in the data of -0.029, suggesting that typical pricing on the platform is highly inelastic. This is consistent with seller pricing incorporating intertemporal demand linkage through search ranking. The finding of inelastic prices may also be influenced by the presence of free products on the marketplace and motivates the inclusion of outside payoff in the supply-side model. I find an age coefficient of -0.00083, suggesting that buyers prefer newer products on the market to old. While 144 category intercepts are not reported in the table, they range from to -0.406 to 0.861. Positive coefficients on keyword length and description length likely reflect differences in quality, complexity, or effort which influence buyer purchase decisions.

On the supply side, I begin by inspecting the choice specific value functions calculated through value function iteration. A sample of surfaces from choice specific value functions can be seen in Figures 4 and 5. Figure 4 holds the rank score fixed and plots choice specific values over the grid of possible price choices and age states. Similarly, Figure 5 holds age fixed while plotting over price choices and ranking score. A few patterns are immediately apparent: optimal price is increasing in current ranking score, consistent with incentives to "invest" in visibility by lowering prices when ranking score is low and then "harvest" when ranking score is high. The optimal price is also slightly decreasing in product age, which is consistent with the typical price skimming pattern observed in other durable goods contexts.

Parameter estimates for the supply model are shown in Table 6. Most sellers on the platform are the high transaction cost type $r = 2$, which reflects the general infrequency of day-level price changes observed in the data. Roughly 16% are low transaction cost type $r = 1$ which incur much smaller penalties to price switching. The α_{rk} parameters capture type-specific outside payoff to pricing and partially rationalize differences between observed price levels and the pricing path suggested by demand estimates and ranking dynamics.

6 Managerial Implications and Counterfactual Simulation

The model suggests a number of key managerial implications. The most notable of which is that sales-ranked visibility generates downward price pressure and especially for new products on the marketplace. Sellers of new products face a low product-specific ranking score due to a lack of historical sales. As seen by the choice specific value function surfaces, there is a strong incentive to discount for low ranking score states. Initial discounts allow a product to generate initial sales and invest in visibility on the marketplace. The returns to visibility can be harvested in later periods by setting a price closer to those implied by demand estimates. Even for mature products, visibility incentives generate downward pricing pressure. This can be seen by taking a derivative of future period demand with respect to current period price. Lower prices today generate increased current period sales in expectation through downward sloping demand. This additional sales quantity increases ranking score which subsequently raises future demand. Thus, a forward-looking profit maximizing seller will always set prices lower than the myopic optimal implied by demand estimates since sellers overvalue sales quantity due to search ranking visibility dynamics. This leads to inefficient pricing on the marketplace: sellers consistently price in the inelastic region of the demand curve and especially implement a poor price discrimination strategy through introductory discounts, as early buyers with highest valuations face the lowest prices.

Introductory pricing strategies by sellers reduce the marketplace operator's revenue. Since the marketplace operator profits by taking a fixed percentage of each transaction, the operator implicitly shares the "costs" of sellers' introductory discounts in attempt to rise in the search rankings. Total visibility on the marketplace is fixed and competition for search ranking position is zero-sum, leading competition amongst sellers to shift buyer attention from one product to another without creating expansion. The marketplace operator thus faces a dilemma similar to that of a wholesaler-retailer channel. Seller agency and competition between sellers - in this case through visibility rather than direct price competition - undermines the potential channel profits that could be realized.

These insights motivate the construction of a counterfactual simulation. In particular, the impact of the marketplace sales-ranked visibility policy could be measured against another policy that does not induce sellers to overvalue sales quantity. I propose revenue-ranked visibility as a relevant comparison. Seller and marketplace operator incentives align under revenue-ranked visibility, as the seller's best price to rise in search rankings is also the profit maximizing price implied by demand and thus the marketplace's preferred price.

In most cases, counterfactual simulation could proceed naturally with draws from the model's implied choice probabilities and updating the relevant state space accordingly. This model, however, introduces a notable complication. The total amount of buyer impressions is both fixed and allocated on the ordinal transformation of ranking scores rather than absolute level. Therefore, any change in sales outcomes on the marketplace implies a new distribution of ranking scores, which implies a new optimal behavior for full information sellers. I resolve this issue through alternatingly drawing behavior based on the current guess of the empirical distribution of ranking scores and updating the empirical distribution based on the most recently drawn behavior. Thus, counterfactual simulation proceeds as follows until convergence:

Marketplace Outcomes Step

1. Draw types for each seller in accordance with supply side estimates.
2. For each product in each period, simulate price choice as a draw from the implied choice probabilities of the choice-specific value function at the given state.
3. Draw the corresponding demand realization for the given price and update states accordingly.
4. Repeat for all products in all time periods.

Update Ranking Score Empirical Distribution Step

1. Collect all ranking scores from the Marketplace Outcomes Step, update the ranking score empirical distribution function $\hat{F}^k(x) = 1/N \cdot \sum_{i=1}^{N} 1_{r_{ik}} \leq x$ in accordance with realized values.

Results of the counterfactual simulation can be seen in Figure 6. From the plotted average price lines, the change from sales-ranked visibility to revenue-ranked visibility eliminates the introductory price discounting incentive and dramatically increases prices for new products on the marketplace. Further, the average price level is slightly higher for mature products, as sellers lose the incentive to lower prices due to future demand considerations. The change to revenue-ranked visibility benefits sellers and the marketplace operator: aggregate seller revenue and therefore operator revenue increase 13% with the policy change. Both purchase incidence and consumer surplus drop by 8% and 14%, respectively.

The potential change to revenue-ranked visibility carries additional implications. Notably, revenue-ranked visibility exacerbates the effects of early sales on long-term product success. Products which sell well initially reach a wider audience, generating more revenue and gaining more exposure, while the opposite is true for products which fail to materialize early sales. Similarly, new products have more difficulty in becoming established and competing with existing products. Sellers previously had ability to take corrective action and rebound from poor sales outcomes by offering discounts in attempt to recover ranking score, but this is not possible under revenue-ranked visibility. Further, sellers which prefer to sell free or discounted products fare much worse under revenue-ranked visibility. A seller of free products would never generate revenue and thus perpetually rank low under the new policy. Lastly, it is unclear if aggregate seller revenue fully captures the marketplace operator's objective. While individual products on the marketplace seem unlikely to carry network effects, the technological platform associated with the marketplace likely does face network effects. The operating firm likely benefits from the creation and utilization of free software available through the marketplace.

7 Conclusions and Future Research

This paper studies visibility policy, seller incentives, and pricing dynamics in the context of a digital goods marketplace. Online marketplaces face the complex task of designing search ranking and other display mechanisms to allocate buyer attention to listed products. Previous literature and industry practice have focused on buyer-side considerations in the design of these systems, often at the cost of neglecting endogenous seller response to incentives introduced by visibility policy. A commonly used visibility policy in app store-like contexts grants additional visibility to products based on sales in previous periods. I study implications of this policy in partnership with the operator of an online platform-marketplace for middleware. This context allows for the rare opportunity to isolate and accurately model the mechanisms that create dynamic incentives in seller pricing. I document that sellers on the marketplace undergo introductory pricing and even price pulsing strategies. These behaviors are rationalized in a dynamic structural model of search ranking visibility on the marketplace. Through counterfactual simulation, I quantify the effect of competition for visibility on aggregate seller revenue and marketplace operator revenue. I propose an alternative search ranking algorithm that the marketplace might pursue which improves the alignment of seller incentives by ranking products on revenue rather than sales quantity. The change from sales-ranked to revenue-ranked visibility increases marketplace revenues by 13% while raising aggregate price levels, reducing purchase incidence by 8%, reducing consumer surplus by 14%, strengthening the effect of initial sales on long-run product success, and potentially alienating free sellers.

Future research plans to evaluate additional counterfactuals and extend the proposed model. Notably, the marketplace operator can mitigate initial discounting and increase revenues by selling sponsored search results through an auction mechanism, creating an alternative mechanism for new product discovery that also generates benefits the operating firm. Evaluating this counterfactual would require specifying a game structure, bidding behavior, and the allocation of visibility associated with auction outcomes. A model extension would look to endogenize the content creation process itself. Sellers would have beliefs about the profitability of selling a product on the marketplace prior to undergoing a costly new product creation and launch process. These beliefs would be informed by prior knowledge and updated through realized outcomes after product release. Then a seller faced with negative experiences may churn out of the marketplace entirely, and the availablility of content on the marketplace becomes influenced by visibility policy. This research fits into a larger agenda around online content creation and how elements of online design can influence user behavior.

Figures and Tables

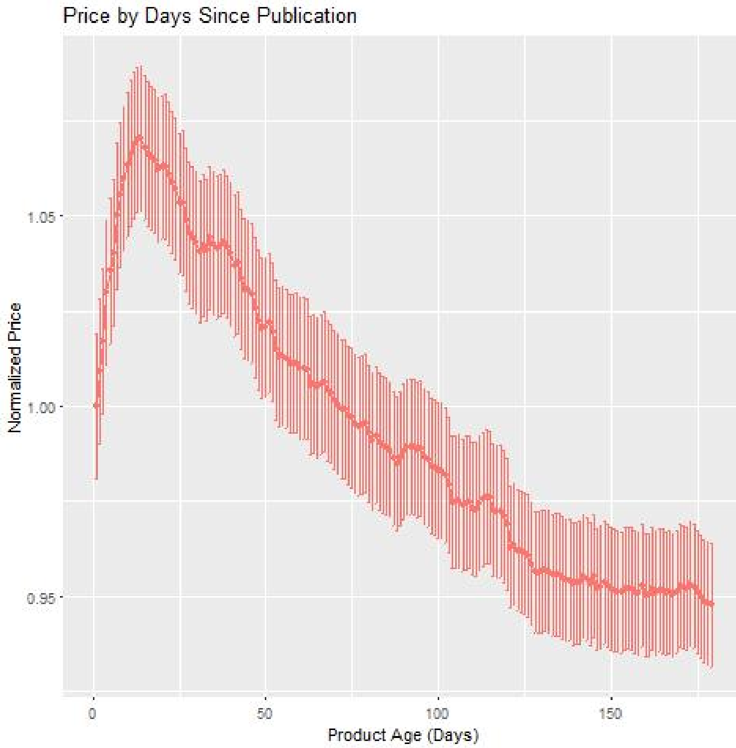

Figure 1: Average Price Path for Products on the Marketplace. Shows the mean price level of products listed on the marketplace of a given age. Error bars represent a 95% confidence interval for the mean price level. Prices are seen to increase by approx. 6% for the first 15 days after product launch and then steadily decline towards a long-run level approx. 5% lower than the starting price level. For confidentiality purposes, the starting price level has been normalized to 1.

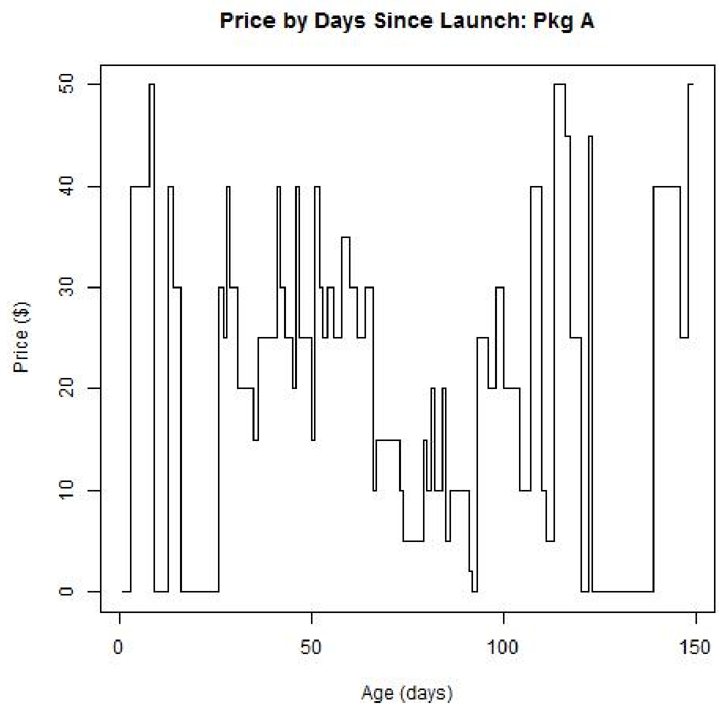

Figure 2: Pricing Trajectory of a Selected Seller. Shows the price level of on a given day of a single product on the marketplace. The seller is seen to alternate between a higher price level and \$0 over periods of a few days. Similar price pulsing strategies are pursued by a subset of sellers on the platform. Price pulsing including \$0 prices seems difficult to justify for a durable good, except in the presence of ranking effects where low prices can "buy" future visibility.

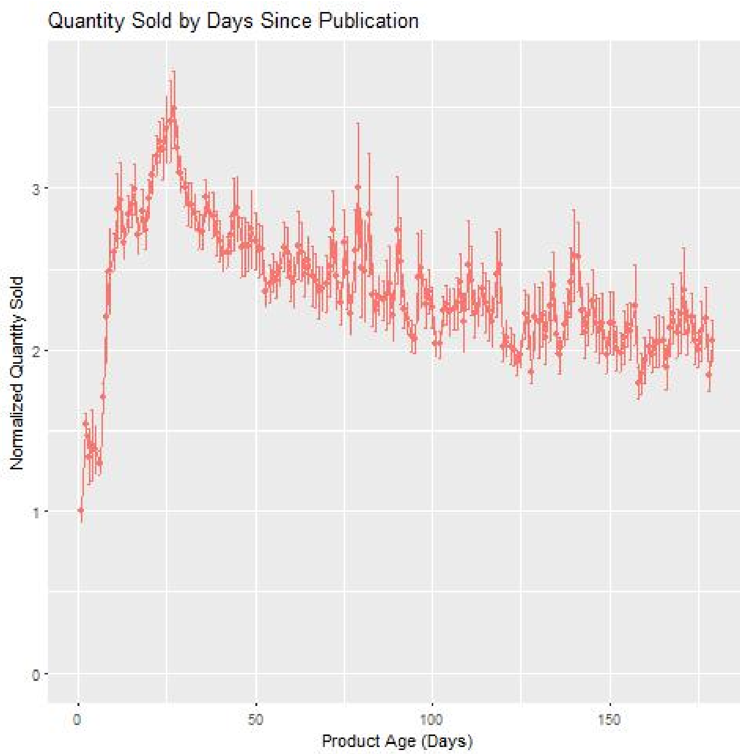

Figure 3: Average Sales Path. Shows the mean quantity sold for products listed on the marketplace of a given age. Error bars represent a 95% confidence interval for the mean quantity sold. Sales are seen to roughly triple over the first 25 days after product launch and decline. For confidentiality purposes, the starting sales level has been normalized to 1.

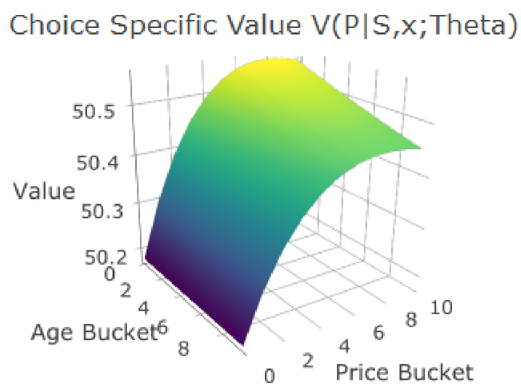

Figure 4: Surface Plot of Choice Specific Value Function by Age and Price. This plot displays choice specific value for a seller as a function of the control variable, price, and the product age state variable. Switching cost is omitted and ranking score is held constant for this visualization. Value is seen to be decreasing in age, and optimal price declines slightly with increased product age from the 8th bucket for newly launched products to the 6th for oldest products.

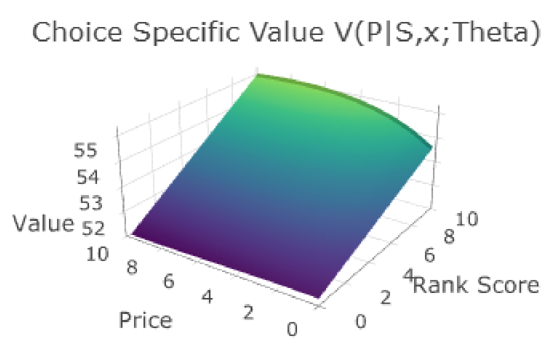

Figure 5: Surface Plot of Choice Specific Value Function by Rank Score and Price. This plot displays choice specific value for a seller as a function of the control variable, price, and the rank score state variable. Switching cost is omitted and product age is held constant for this visualization. Predicted optimal price choice is dependent on rank score state - at the lowest rank score the 2nd price level is optimal while at the highest rank score the 8th price level is optimal.

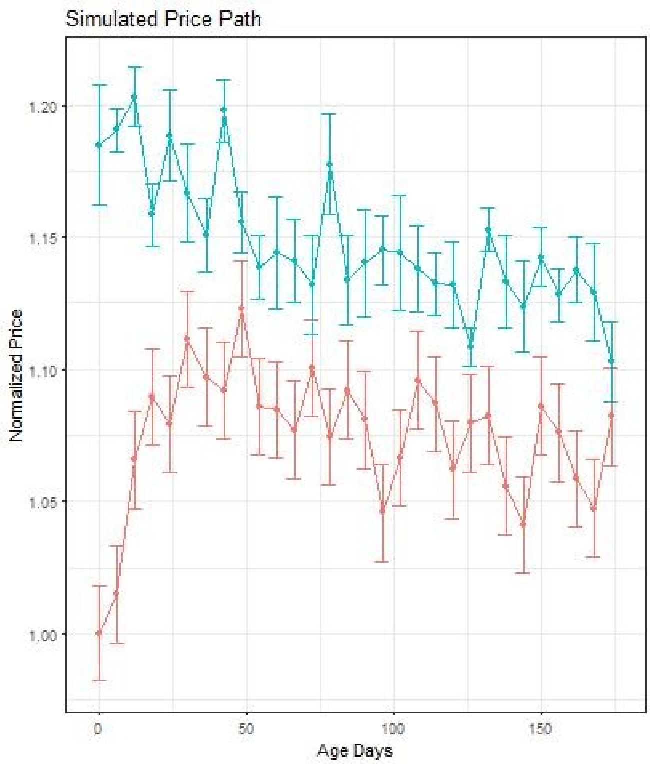

Figure 6: Counterfactual Simulation: Revenue-Ranked Visibility. This plot displays simulated mean product prices by product age under the current sales-ranked visibility policy and a proposed revenue-ranked visibility policy. Prices are notably higher under revenue-ranked visibility and lack the initial price discount seen under simulation of the current sales-ranked policy.

Table 1: Normalized Price by Age

	Dependent variable:	
	Price	
	Age <= 15	Age > 15
Age	0.002***	−0.0002***
	(0.0002)	(0.00003)
Product Fixed Effects	Yes	Yes
R^2	0.989	0.970
Note:	*p<0.1; **p<0.05; ***p<0.01	

Table 2: Normalized Sales by Age

	Dependent variable:	
	Sales Quantity	
	Age <= 25	Age > 25
Age	0.092***	−0.005***
	(0.007)	(0.001)
Age^2		
Product Fixed Effects	Yes	Yes
R^2	0.182	0.071
Note:	*p<0.1; **p<0.05; ***p<0.01	

Table 3: Regressions of Normalized Sales on Normalized Price

	Dependent variable:		
	(normalized) Sales Quantity		
(normalized) Price	0.174***	−0.043***	−0.030**
	(0.006)	(0.013)	(0.012)
Constant	2.221***		
	(0.014)		
Product Fixed Effects	No	Yes	Yes
Ranking Score Controls	No	No	Yes
R^2	0.0003	0.066	0.160
Note:	*p<0.1; **p<0.05; ***p<0.01		

Table 4: Product Impressions vs Previous Period Sales

	Dependent variable:
	Product Impressions
Previous Period Sales Quantity	1.236***
	(0.282)
Lagged Impressions	0.484***
	(0.014)
Product Fixed Effects	Yes
R^2	0.559
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 5: Demand Model Estimates

	Dependent variable:
	$log(s_{j1t}/s_{j0t})$
price	-0.00142***
	(0.00039)
age_days	-0.00083***
	(0.00034)
keywords_wordcount	0.00779***
	(0.00018)
description_wordcount	0.00089***
	(0.00015)
Category Intercepts	Yes
R^2	0.0326
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 6: Supply Model Estimates

Variable	*Parameter*	*Std Err*
Fraction Type 1 (ρ)	0.1582	0.0609
Transaction Cost 1 (λ_1)	0.3188	0.1098
Transaction Cost 2 (λ_2)	4.022	0.2402
Outside Payoffs:		
$\alpha_{1,1}$	0.2951	0.2160
$\alpha_{1,2}$	0.2377	0.1354
$\alpha_{1,3}$	-0.1488	0.2308
$\alpha_{1,4}$	-0.1946	0.2759
...		
$\alpha_{2,16}$	0.1691	0.0987
$\alpha_{2,17}$	0.1183	0.2304
$\alpha_{2,18}$	-0.1902	0.1554
$\alpha_{2,19}$	-0.2574	0.1062
$\alpha_{2,20}$	-0.3012	0.1094

References

Nick Arnosti, Ramesh Johari, and Yash Kanoria. Managing congestion in dynamic matching markets. *ACM Conference on Economics and Computation*, 2014.

Kyle Bagwell. Introductory price as a signal of cost in a model of repeat business. *The review of economic studies*, 1987.

Subramanian Balachander and Kannan Srinivasan. Modifying customer expectations of price decreases for a durable product. *Management science*, 1998.

Steven Berry, James Levinsohn, and Ariel Pakes. Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society*, pages 841–890, 1995.

Timothy Bresnahan, Xing Li, and Pai-Ling Yin. Paying incumbents and customers to enter an industry: Buying downloads. *Working Paper*, 2015.

Luis Cabral and Ali Hortacsu. The dynamics of seller reputation: Evidence from ebay. *The Journal of Industrial Economics*, 58.1:54–78, 2010.

Gerard P Cachon. Supply chain coordination with contracts. *Handbooks in operations research and management science*, 2003.

Olivier Chapelle and Ya Zhang. A dynamic bayesian network click model for web search ranking. *Proceedings of the 18th international conference on World wide web*, 2009.

Yuxin Chen and Song Yao. Sequential search with refinement: Model and application with clickstream data. *Management Science*, 2016.

Hana Choi and Carl Mela. Online marketplace advertising. *Working Paper*, 2016.

Ronald H. Coase. Durability and monopoly. *The Journal of Law and Economics*, 1972.

John Conlisk, Eitan Gerstner, and Joel Sobel. Cyclic pricing by a durable goods monopolist. *The Quarterly Journal of Economics*, 1984.

Jean-Pierre Dube, G. Hitsch, and Pranav Jindal. Estimating durable goods adoption decisions from stated preference data. *Working Paper, University of Chicago*, 2009.

Joseph Farrell and Garth Saloner. Installed base and compatibility: Innovation, product preannouncements, and predation. *The American economic review*, 1986.

Andrey Fradkin. Search frictions and the design of online marketplaces. *NBER Working Paper*, 2014.

Matthew Goldman and Justin Rao. Experiments as instruments: heterogeneous position effects in sponsored search auctions. *http://dx.doi.org/10.2139/ssrn.2524688*, 2014.

Justin Huang and Sridhar Narayanan. Effect of feature recognition on content creation and engagement: Experimental results from a social network. *Working Paper*, 2016.

Thorsten Joachims. Optimizing search engines using clickthrough data. *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, 2002.

Michael Katz and Carl Shapiro. Product introduction with network externalities. *The Journal of Industrial Economics*, 1992.

Paul Milgrom and John Roberts. Informational asymmetries, strategic behavior, and industrial organization. *The American Economic Review*, 1987.

K. Sridhar Moorthy. Product and price competition in a duopoly. *Marketing Science*, 1988.

Harikesh Nair. Intertemporal price discrimination with forward-looking consumers: Application to the us market for console video-games. *Quantitative Marketing and Economics*, 5.3:239–292, 2007.

Chakravarthi Narasimhan. Incorporating consumer price expectations in diffusion models. *Marketing Science*, 1989.

Sridhar Narayanan and Kirthi Kalyanam. Position effects in search advertising and their moderators: A regression discontinuity approach. *Marketing Science*, 2015.

Chris Nosko and Steven Tadelis. The limits of reputation in platform markets: An empirical analysis and field experiment. *National Bureau of Economic Research*, 2015.

Jean-Charles Rochet and Jean Tirole. Platform competition in two-sided markets. *Journal of the european economic association*, 1.4:990–1029, 2003.

John Rust. Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica*, 55(5):999–1033, 1987.

Steven C Salop. Monopolistic competition with outside goods. *The Bell Journal of Economics*, 1979.

Inseong Song and Pradeep Chintagunta. A micromodel of new product adoption with heterogeneous and forward-looking consumers: Application to the digital camera category. *Quantitative Marketing and Economics*, 1.4:371–407, 2003.

Nancy L. Stokey. Intertemporal price discrimination. *The Quarterly Journal of Economics*, 1979.

Monic Sun and Feng Zhu. Ad revenue and content commercialization: Evidence from blogs. *Management Science*, 59(10):2314–2331, 2013.

Steven Tadelis and Florian Zettelmeyer. Information disclosure as a matching mechanism: Theory and evidence from a field experiment. *American Economic Review*, 105(2):886–905, 2015.

Olivier Toubia and Andrew Stephen. Intrinsic vs. image-related utility in social media: Why do people contribute content to twitter? *Marketing Science*, 32.3:368–392, 2013.

Hal R Varian. A model of sales. *The American Economic Review*, 1980.

Glen Weyl. A price theory of multi-sided platforms. *The American Economic Review*, 100.4: 1642–1672, 2010.

Song Yao and Carl Mela. A dynamic model of sponsored search advertising. *Marketing Science*, 2011.

Xiaoquan Zhang and Feng Zhu. Group size and incentives to contribute: A natural experiment at chinese wikipedia. *The American Economic Review*, (101.4):1601–1615, 2011.