PUBLISH ON DEMAND

Content On The Go: The Economics of the Market for Mobile Apps

FALL 2013

UPDATE

mediaX
STANFORD UNIVERSITY

mediaX
at Stanford University

mediaX connects businesses with Stanford University's world-renowned faculty to study new ways for people and technology to intersect.

We are the industry-affiliate program to Stanford's H-STAR Institute. We help our members explore how the thoughtful use of technology can impact a range of fields, from entertainment to learning to commerce. Together, we're researching innovative ways for people to collaborate, communicate and interact with the information, products, and industries of tomorrow.

Content on the go: The economics of the market for mobile apps

Bar Ifrach Ramesh Johari*

September 11, 2013

Media X Final Report

Abstract

The growth of mobile applications on smartphones and tablets ("apps") ranks as one of the most astonishing technological developments in recent past. Over 700,000 apps are available for immediate download from app markets (e.g., App Store and Google Play). These marketplaces are a significant disruptive change in the way content is created and consumed. On the supply side, they provide content creators direct, instantaneous, and popular distribution systems where they can implement their own marketing and pricing policies, cutting out middlemen.

Taking a combined data-driven and structural analysis approach, this paper focuses on the relationship between pricing decisions and marketplace visibility. Our aim is to empower content creators by offering strategic guidance on how to leverage the marketplaces' flexibility. Specifically, the market platforms feature "top-ranked" charts that list apps by number of downloads. A high position in these charts is followed by a remarkable boost in demand, according to industry sources. We call the effect of top-rank position on future sales an indirect effect. First, we postulate a reduced form model to estimate the magnitude of this indirect effect. Our results show that it is statistically significant and substantial. Second, we study app pricing decisions in a model that incorporates our earlier findings. Surprisingly, we find that accounting for the indirect effect may give rise to optimal price cycles, where the seller alternates between a high price to boost revenue and a low one to enhance visibility. We find numerous evidence supporting this pricing behavior in practice.

*Both at Management Science & Engineering, Stanford University. {bifrach, ramesh.johari}@stanford.edu. The authors thank Media X and Konica Minolta for their generous support.

1 Introduction

The growth of mobile applications on smartphones and tablets ("apps") ranks as one of the most astonishing technological developments in recent past. Over 700,000 apps, either free or paid, are available for immediate download from designated app markets (e.g., App Store and Google Play), generating over $20 million in daily revenue that is rapidly growing[1]. These app marketplaces, the earliest launched in July 2008, are a significant disruptive change in the way content is created and consumed, grabbing consumers' attention away from legacy media, e.g., print, as well as modern media, e.g., web browsing on a PC.

On the demand side, the marketplaces offer users rich content utilizing the functionality of their mobile devices. In particular, by using the unique dimensions of user experience available on modern mobile devises: location, image and touch based interfaces. This is a whole new field of functionality and content consumption that is rapidly replacing more traditional mediums. For example, sales of digital cameras has dropped by over 40% in 2013 as smartphones replace digital cameras as the latter did to to film cameras a decade ago[2].

On the supply side, these platforms provide content creators direct, instantaneous, and highly popular distribution systems where they can implement their own marketing and pricing policies, cutting out middlemen. However, for a content creator, making sensible business decisions requires an understanding of the economics underlying this market—including competition, features of the market platform, and pricing.

Taking a combined data-driven and structural analysis approach, this paper studies various aspects of the relationship between pricing decisions and marketplace visibility. Our aim is to empower individual content creators by offering strategic guidance on how to leverage the marketplaces' flexibility.

We focus in particular on the role of rankings and recommendations in driving demand. The market platforms offer a number of recommendation systems designed to harness the so called "wisdom of the crowds" to help users choose what to download in the plethora of apps. The most salient among them are the "top-ranked" charts that list apps by number of downloads, as well as some secondary popularity indicators. A high position in these charts is followed by a remarkable boost in demand, according to both industry sources (Surikate and GfK, 2012) and empirical research (Carare, 2012). We call the effect of top-rank position on future sales an *indirect effect* or *visibility effect*, to distinguish it from the direct relationship between the past sales and rank (since rank is a measure of past sales in comparison to those of competing apps).

We proceed in two directions. First, we postulate a reduced form model to estimate the magni-

[1] See Distimo (2013).
[2] See Wakabayashi (2013).

tude of the indirect effect of ranks on sales, employing time series of top-ranked charts from June 2012 to March 2013. This model isolates the indirect effect, as we outline in more detail below. Our results show that the indirect effect is statistically significant and substantial for the top-paid ranking chart. In particular, we find a significant visibility effect for apps in positions 1-20, but not to apps in positions 21 and forward. Compared with the app in rank position 20, the app in rank position 1 make about 90% more downloads due to the visibility effect. However, the total difference in downloads between the two apps is likely to be much higher because it includes the direct effect as well.

With this effect in hand, we study app pricing decisions in a stylized model that incorporates our earlier findings. Surprisingly, we find that accounting for the indirect effect may give rise to optimal *price cycles*, where the seller alternates periodically between a high price and a low one to boost revenue in the first and market visibility in the latter. We find evidence in the data supporting this pricing behavior in practice.

The reminder of the paper is organized as follows. Subsection 1.1 describes related literature. Section 2 describes the top-ranked charts and our dataset and Section 3 describes the estimation model and results. Section 4 describes the pricing problem and pricing cycles and Section 5 concludes.

1.1 Literature Review

This work is related to a number of papers that can divided to two groups. The first group holds two papers that study the marketplace for apps directly. Garg and Telang (2012) infer the relationship between rank position and downloads in App Store, assuming the download distribution follows a power law and utilizing an equation relating download and grossing ranks. We make a similar assumption on the distribution of downloads and use heir estimates to identify the visibility effect.

Carare (2012) has a similar objective to our paper, but takes a different approach.Carare (2012) estimates the dollar value of visibility for each top-ranked position, i.e., how much bigger is users' willingness-to-pay for a top-ranked app compared with an unranked one. We, instead, estimate the percentage increase in downloads resulting from an increase in rank position. Since estimating the monetary equivalent of a rank position, Carare (2012) makes use of price variations that are not very common in the marketplace and, as a result, requires a strong assumption that all apps share the same price elasticity. We take a different approach that makes use of rank variations instead of price variation. Because of the different estimation model, comparing the estimation models requires imposing a common price elasticity assumption, which we do not make.

Carare (2012) is also part of the second group of papers related to ours that investigate the visibility effect of top-ranked positions on demand. The main challenge to this empirical question is disentangling causality from correlation between rank positions and demand, since rank positions depend on relative demand levels. Our approach in tackling this endogeneity of explanatory variables resembles that in Sorensen (2007) that estimatea the visibility effect of the *New York Times* top-ranked books list. Sorensen (2007) finds instances constituting rank shifters that are not correlated with variations in demand, due to mistakes and delays in the release of information. Unlike with top-ranked books on the New York Times that are released weekly, top-ranked charts in App Store are updated multiple times a day and are completely automated. This, however, opens a new opportunity to finding rank shifters in the forms of no-swaps, as we explain in detail in Section 3.

2 Top Ranked Charts

Ranking lists capture the popularity of different apps based on their recent demand, rewarding popular apps with salient market visibility and a trendy appeal. However, as much as ranking lists reflect underlying demand patterns, they may also set them by boosting the demand of already top-ranked apps over less popular ones, further skewing the demand distribution. If such indirect feedback exists, we can think of the demand for an app as a function of its attributes (e.g., functionality and graphics), its price and its rank position.

This indirect effect of rankings on demand, if exits, has important implications for both app developers and for the marketplace operator. From the developers' perspective, this effect should be taken into account when designing pricing and marketing strategies. A strong effect of ranking position on demand should also concern the marketplace operator seeking to support discovery of new apps, as incumbent apps may have an advantage over newly launched ones that perhaps offer a better user experience.

In the reminder of this section we discuss various aspects of the top-ranked charts and define the main notation for the paper.

Let I_t be the set of all apps available at time t. We denote by d_{it} the number of downloads of app $i \in I_t$ at time t, and analogously denote by s_{it} the sales (revenue) that include both the download price and all revenue from in-app purchases. The time elapsing from t to $t+1$ varies from 6-8 hours to a day and is stated when relevant.

At each time t the marketplace operator (e.g., Apple's App Store and Google Play) publishes a number of top-ranked lists. A ranking list l at time t is a partially ordered list that grades apps in a subset $A_t^l \subset I_t$ by some underlying function $h^l(S_t, D_t)$, where $D_t := (d_0, \ldots, d_t)$ and

$S_t := (s_0, \ldots, s_t)$. Let $r_{it}^l \in \{1, 2, \ldots, K\} \cup \{\text{NR}\}$ be the l rank of app i at time t, where NR stands for *not ranked*. Then we assume that the app in rank $k \leq K$ is the k-th largest measured by h^l, and apps that are smaller than the K-th largest app are not ranked. Formally, denote by

$$N_t^l(i) := \#\{j \in A_t^l | h^l(S_{jt}, D_{jt}) > h^l(S_{it}, D_{it})\} + 1$$

the number of apps with $h(S_{tj}, D_{tj})$ larger than that of app i. We have

$$r_{it}^l := \begin{cases} N_t^l(i), & \text{if } N_t^l(i) \leq K \\ \text{NR}, & \text{otherwise.} \end{cases} \tag{1}$$

The most visible ranking lists published in both AS and GP are the top-paid, top-free and top-grossing rankings over all apps. The first two capture download popularity for apps with a positive download price and for free apps, respectively. The grossing ranking captures the revenue generated by apps, both from the download price and from in-app purchases, and includes both free and paid apps. Both marketplaces similarly include top-free and top-paid rankings for each app category such as games and productivity, where only AS publishes grossing ranks per category. In the reminded of the paper we will index ranking lists by the category they cover and by the list type: 'f' for top-free, 'p' for top-paid and 'g' for top-grossing, e.g., $r_t^{\text{overall},f}(k)$ is the k-th most popular free app over all categories. A number additional rankings considering new, trending and tablet oriented apps are available.

For example, on September 10, 2013 *The Sims 3*, a simulation game app, was in the 173 and 299 positions in the overall top-paid and top-grossing charts, respectively. Thus, we write $r^{\text{overall,p}} = 173$ and $r^{\text{overall,g}} = 299$, where we omitted the time and app index for clarity. In addition, the app is in positions 73 and 212 apps for top-paid and top-grossing ames, respectively, with the notation $r^{\text{games,p}} = 73$ and $r^{\text{games,g}} = 212$. The overall paid chart considers all apps $A_t^{\text{overall,p}} = I_t \cap \{\text{app}$ i's price is positive at $t\}$, while the games chart is a subset of the former including games only. Hence, if an app in top-ranked in the overall chart, its category position will be lower, as see in this example.

The lower the rank position of an app, the more popular it is. In spite of that, when we write that an app A is more highly ranked than app B, we mean that A is more popular than B and with a lower rank position. Similarly, when we write a that an app is highly ranked we mean that its rank position is low.

2.1 Ranking algorithm

The platforms do not disclose the algorithm measuring the popularity and sales of apps used to construct the ranking lists, i.e., the functions h^l. The common belief in the industry is that this algorithm is based on a weighted average of the last 3-4 days of downloads for ranks top-free and top-paid lists or of sales for top-grossing[3]. This input is likely to be supplemented with other popularity metrics such as click-through-ratios of the apps icon, consumer reviews, and social media popularity.

Our main dataset does not allow us to reverse engineer the ranking algorithm, since it includes apps' ranking data, but not their download data (see Subsection 2.2). Nevertheless, a small number of app developers shared with us their download data.

Figures 1 and 2 plot two time series of daily rank positions and downloads for two apps, app C and app F in the *Catalog* and *Food and Drinks* categories, respectively. The rank positions are for the chart of top-paid apps in those respective categories. The inverse relationship between the two variables is clear: the higher downloads the lower the rank. It is particularly clear during peaks and drops in download levels[4].

Thus motivated, Appendix B develops an estimation procedure to explore the principal factors controlling h^l. We find that downloads explain 30-58% of variations in rank positions, a figure that we find to be quite high, since rank positions are also strongly affected by downloads for other apps. In addition, and in contrast to the older industry findings reported above, we find that lagged downloads have little to no effect on ranks. Thus, we assume throughout the paper that the function $h^l(D_{it}, S_{it}) = d_{it}$ whenever l is a popularity chart and $h^l(D_{it}, S_{it}) = s_{it}$ whenever l is a grossing chart.

2.2 Data

Our dataset spans all ranking charts—free, paid, grossing for the overall charts and all category charts—from June 2012 to March 2013. The rank positions are recorded 3-4 times daily and includes the price for paid apps. In addition, we have less regular supplementary data including each app's category, version, size in megabytes, and other content available on the app's iTunes page. Ranking charts include the top 400 positions ($K = 400$).

[3]See faberNovel (2010).

[4]In addition, we see periodic fluctuations in demand, especially for app F, due to a weekend effect.

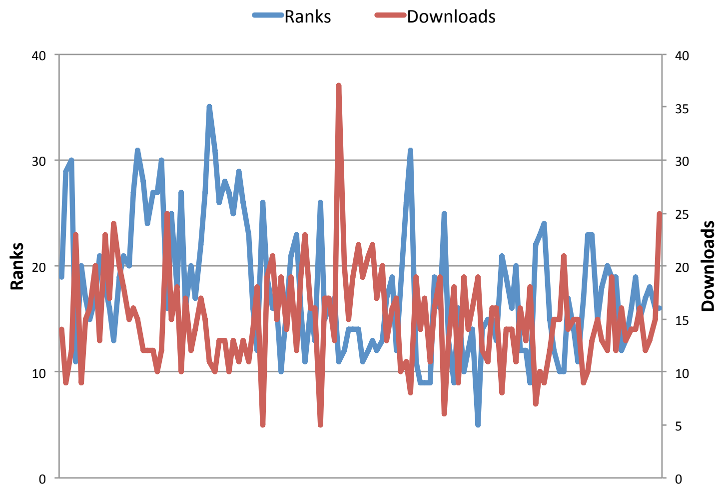

Figure 1: Time series of daily ranks and downloads for app C (126 days in early 2013).

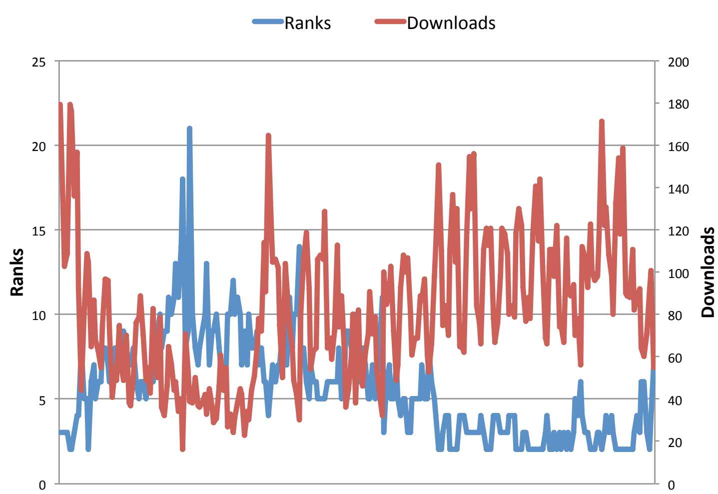

Figure 2: Time series of daily ranks and downloads for app F (252 day stating in end of 2012).

3 Marketplace Visibility and Demand

In this section we estimate the impact of top-charts visibility on sales. To focus the discussion, suppose that downloads of app i at time t follow a demand function

$$d_{it} = \exp(\beta_i' X_{it} + f(r_{it}) + \delta_t + \epsilon_{it}), \tag{2}$$

where X_{it} are app specific attributes (e.g., category and price) with coefficients β_i (e.g., price elasticity), f captures the indirect effect of the rank position on downloads, δ_t is a market wide disturbance, and ϵ_{it} is an app specific disturbance. This calls for the estimation equation for f, specified in differences for convenience,

$$\Delta\tilde{d}_{it} = \beta_i' \Delta X_{it} + \Delta f(r_{it}) + \Delta\delta_t + \Delta\epsilon_{it}, \tag{3}$$

where for any time series variable $\Delta x_t := x_t - x_{t-1}$ and $\tilde{x} := \log(x)$ for any variable x.

In our data set, $\Delta X_{it} = 0$ in all cases, except for price changes that are not common, and would not allow for estimation of app specific price elasticity. Therefore, we omit observations before and after a price a change, and consider the equation

$$\Delta\tilde{d}_{it} = \Delta f(r_{it}) + \Delta\delta_t + \Delta\epsilon_{it}, \tag{4}$$

Two hurdles to the proposed estimation procedure arise in our context. First, variations in downloads levels naturally induce variations in rank positions, resulting in a severe endogeneity problem. Namely, the rank position at time t depends on the downloads at time t, and consequently on the idiosyncratic disturbances, ϵ_{it}. We resolve this problem by exploiting a natural shifter to rank position that is unrelated to variations in demand.

The second hurdle results from the lack of download data (d_{it}), which developers are reluctant to share[5]. To circumvent this problem, we take advantage of the availability of apps in both top-download and top-grossing apps. We use the latter as a proxy for demand under the assumption that ranks and downloads follow a power law relationship as assumed in Carare (2012) and Garg and Telang (2012) in the app market, and in Brynjolfsson, Smith, and Hu (2003) and Chevalier and Mayzlin (2006) for book ranking. In the reminder of this subsection we introduce our approach and present an estimation procedure that produces an unbiased estimate for f.

[5] As we mention above, a number of developers shared with us their downloads information. However, none of the relevant apps is ranked high enough to gain from the indirect effect.

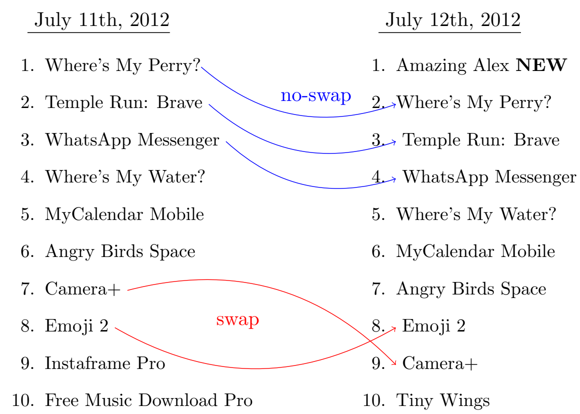

Figure 3: Illustration of swap and no-swap on overall top-paid ranking chart

3.1 Rank Shifters

Ranks change over time for two reasons: (a) reflecting variations in the demand for the app itself, and (b) reflecting variations in the demand for other apps. For example, between July 11^{th} and July 12^{th} 2012, a number of highly ranked apps dropped one position in the overall top-paid list, following the launch of the much anticipated game *Amazing Alex* straight to the top of the chart (the game later flopped). A popular texting app, *WhatsApp Messenger*, dropped from position 3 to 4, suggesting that this decrease in rank is likely unrelated to a change in its underlying demand. On the same dates a photography app, *Camera +*, dropped from position 7 to 9, while the rank of an emoticon app, *Emoji 2*, jumped up from rank 9 to 8, see Figure 3 for an illustration. We call this a *swap*: the ranking order of two apps on the same list changed. A swap indicates an underlying change in demand of at least one of the swapping apps: either the demand of the lower ranked app at time t increases or that of the higher ranked app decreases (or both).

We use this distinction to isolate the effect of ranks on demand using our dataset that includes a time series of ranking data and prices. Namely, we identify a subset of app indexes and time points, W with typical element (i, t) such that the orthogonality condition

$$\mathsf{E}[\Delta f(r_{it})\Delta\epsilon_{it}] = 0 \text{ for } (i,t) \in W \tag{5}$$

holds. This approach is made feasible with the quick update of the top-ranked charts that result in many rank changes.

Denote by $a_t(r) = \{i | r_{it} = r\}$ the app at rank r at time t on some chart l, whose index is suppressed for notational clarity. We say that app i did *not* k-swap on chart l at time t if the block of k apps surrounding app i on the chart do not change order on the rank at $t+1$ and remain adjacent. Formally, if $a_t(r_{it}+j) = a_{t+1}(r_{it+1}+j)$ for all $j = -k/2, -k/2+1, \ldots, k/2$. For example,

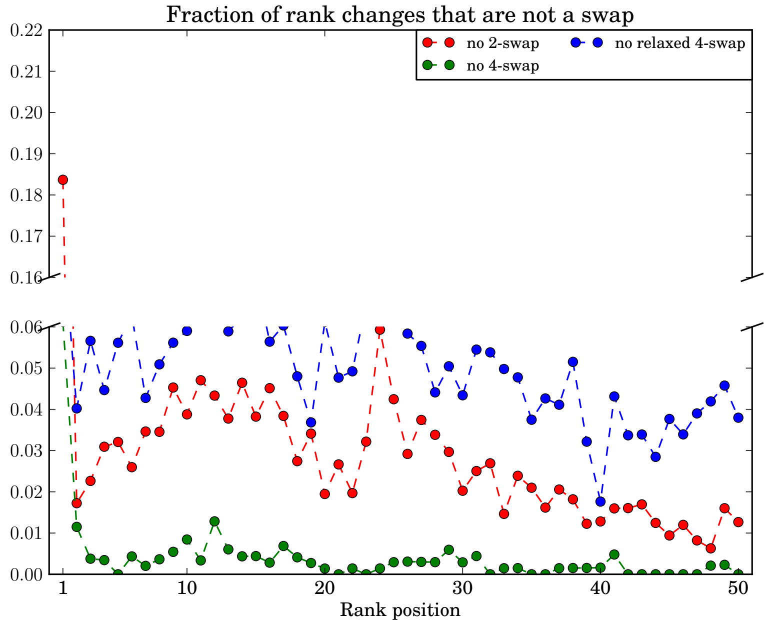

Figure 4: Fraction of rank changes that are no-swap per rank position.

an app that stayed at position 10 at both periods t and $t+1$ did not 4-swap only if the apps at positions 8,9,11 and 12 did not change rank positions from t to $t+1$ as well. Similarly, if the app dropped to rank position 8 that would not be considered a 4-swap only if the apps at position 8-12 would drop 2 positions each as well. Any other case would be called a k-swap. We denote by W^k the set of app indexes and time points (i,t) for each a k-swap did not take place. Note that a single app and single time point can appear in W^k multiple times.

We will consider two sets of no-swaps in our estimation model, W^2 and a modified W^4, denoted by $\widetilde{W}^4$, where the latter includes observations such that if one "removes" one app from the 5 adjacent ones, the remanning apps form a no 4-swap, i.e., the observation is almost a no 4-swap. Formally, $a_t(r_{it}+j) = a_{t+1}(r_{it+1}+j+k_j)$ for $j=-2,-1,0,1,2$, $k_j \in \{0,1\}$ with $k_j \leq k_{j+1}$ and $k_2 \leq 1$. We look at this modified definition, since there are not enough observations in the set W^4 to establish statistical significance, see Figures 4 and 5.

3.2 Cross Chart Variation

The second challenge in applying the estimation procedure described above is the lack of download data. To circumvent this we take advantage of the second time series available to us: the grossing ranks of apps. The grossing chart is similar to the download charts, except that it is based on sales, and not on downloads. In addition, the top-grossing charts are less saliently displayed on the app store, and of lesser interest to consumers (Google Play does not release category top-

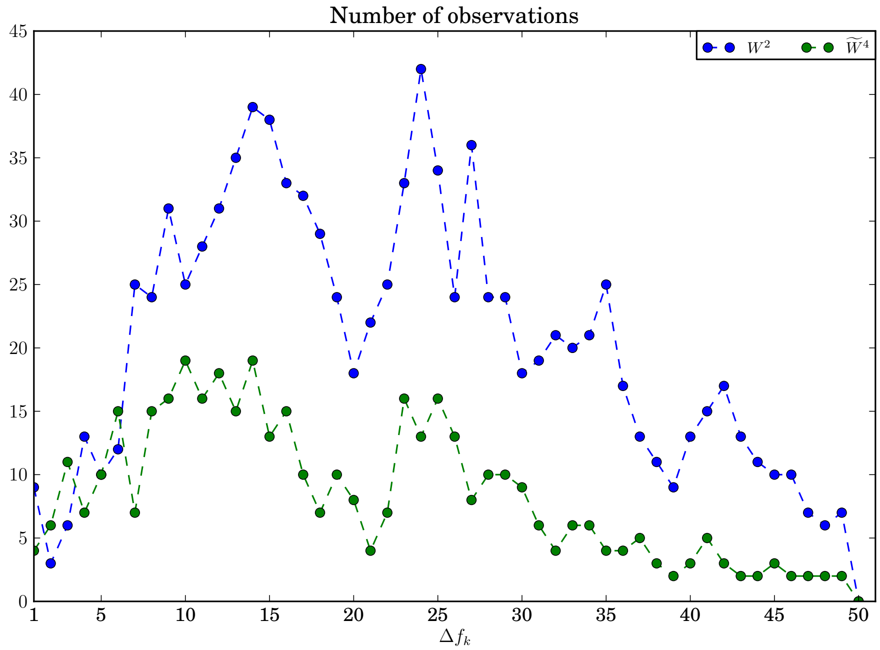

Figure 5: Number of no-swap observations per rank position.

grossing charts) . The revenue that accounts for the top-grossing rank is that made with direct payments by consumers—download price and in-app purchases.

Following Carare (2012) and Garg and Telang (2012), we suppose that the ranks represent a power law of downloads and sales for some parameters A^l and θ^l

$$r_{it}^l \approx A^l (d_{it})^{\theta^l} \text{ for } l \text{ either a free or download chart}$$

and

$$r_{it}^l \approx A^l (s_{it})^{\theta^l} \text{ for } l \text{ a grossing chart.}$$

Sales are a function of download given by $s_{it} = (p_{it} + \nu_{it})d_{it}$, where ν_i is the fraction of revenue attributed to in-app purchases. This gives rise to linear relationship between downloads and grossing ranks

$$\Delta \tilde{d}_{it} = \theta^g \Delta r_{it}^g. \tag{6}$$

With equation 6 we are able to replace the variations in the downloads with variations in grossing ranks and arrive at our main estimation equation

$$\Delta \tilde{r}_{it}^g = (\theta^g)^{-1} \left[\Delta f(r_{it}) + \Delta \delta_t + \Delta \epsilon_{it} \right]. \tag{7}$$

The inverse of the parameter θ^g was estimated by Garg and Telang (2012) to be -1.163 with standard error of 0.011. This will be used to isolate the effect of visibility.

3.3 Estimation results

Denote by $\Delta f_k := f(k) - f(k+1)$, the marginal visibility effect of rank position k. Figures 4 and 5 show the number and fraction of observations that include a variation in Δf_k for $k = 1, \ldots, 50$ for both W^2 and $\widetilde{W}^4$. In addition, the sets W^2 and $\widetilde{W}^4$ include 5561 and 5098 observations, respectively, that did not include a rank change. Ideally we would like to estimate a parameters for each increment in f, Δf_k for $k = 1, 2, \ldots$. However, Figure 5 shows that we do not have enough rank change observations to identify each increment, especially in $\widetilde{W}^4$. Therefore, we assume that f is piecewise linear on bins of 5 adjacent rank positions $[1, \ldots, 5], [6, \ldots, 10], \ldots$. Namely, we would like to estimate the slopes $\beta_{1,5}, \beta_{6,10}, \ldots$ for the corresponding bins, such that $\Delta f_k = \beta_{\mathrm{bin}_k}$, where bin_k is the bin containing rank position k.

Denote by

$$x_{itk} := \begin{cases} \left|[r_{it-1}, r_{it-1}+1, \ldots, r_{it}-1] \cap \mathrm{bin}_k\right| & \text{if } r_{it} > r_{it-1} \\ -\left|[r_{it}, r_{it}+1, \ldots, r_{it-1}-1] \cap \mathrm{bin}_k\right| & \text{if } r_{it} < r_{it-1} \\ 0 & \text{if } r_{it} = r_{it-1}, \end{cases}$$

where $|S|$ is the cardinality of set S. Namely, x_{itk} is the rank increment of app i at time t restricted to bin k. For example, if $r_{it} = 2$ and $r_{it-1} = 3$ then $x_{it1} = -1$ and $x_{itk} = 0$ for all $k = 2, 3, \ldots$. By the same principle, $r_{it} = 7$ and $r_{it-1} = 5$, then $x_{it1} = 1$, $x_{it2} = 1$, and $x_{itk} = 0$ for all $k = 3, 4, \ldots$.

Note that under assumption that f is piecewise linear, $\Delta f_{it} = \sum_{k=1}^{10} \beta_{\mathrm{bin}_k} x_{itk}$. Therefore, our principle estimation equation is

$$\Delta r_{it}^g = \beta_0 + \sum_{k=1}^{10} \beta_{\mathrm{bin}_k} x_{itk} + \xi_{it}. \tag{8}$$

The estimators and standard errors for this principle model under $\widetilde{W}^4$ and W^2 are reported in Tables 1 and 2 in Appendix A. We find that the coefficients for bins greater than $[16, 20]$are not significant under the set $\widetilde{W}^4$ and remove them from the estimation equation sequentially, that is we cannot conclude that the visibility of positions 21 and greater increases downloads. This supports the industry's focus on penetrating to the top 25 top-ranked charts.

Recall from (7) that $\beta_{\mathrm{bin}_k} = (\theta^g)^{-1} \Delta f_l$ for $l \in \mathrm{bin}_k$. To obtain the effect on *downloads*, we first multiple the coefficients by θ^g to obtain estimates for Δf_k. Following (6), we plot in Figure 6 the estimated increase in downloads due to the visibility effect over rank 20, the lowest rank for which

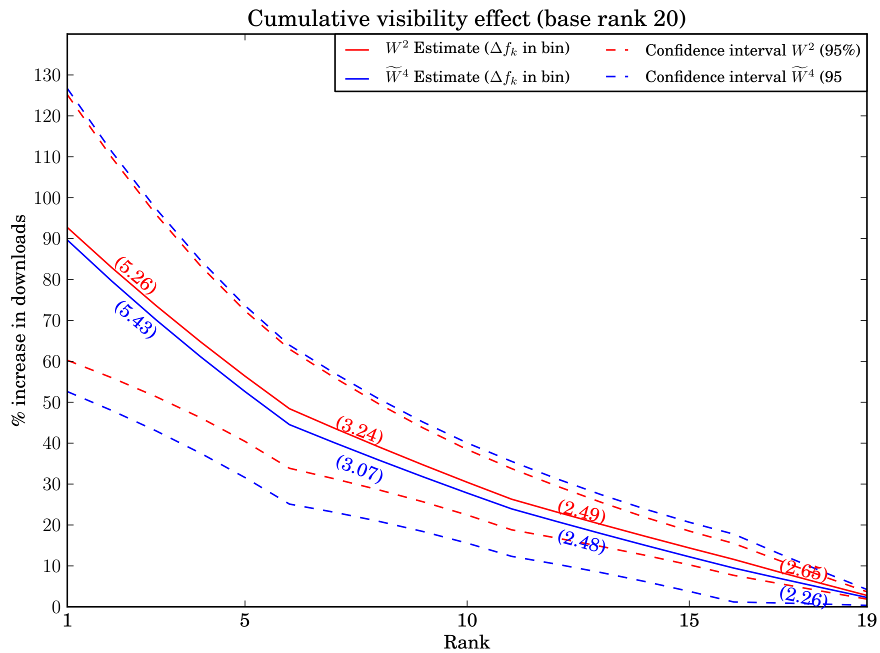

Figure 6: Cumulative visibility effect over rank 20. Assumes $\theta^g = -1.163$.

we identify the effect in the set $\widetilde{W}^4$. In the set W^2 all the coefficients for bins $[1, 5]$ to $[46, 50]$, except $[41, 45]$ are significant. For this we compute the estimate of $\exp(f(n) - f(20)) - 1$ and it variance using the Delta Method.

For set $\widetilde{W}^4$, we estimate that the app in rank 1 gains additional 89% in downloads over than in rank 20 due to the greater visibility of its rank position, and at least 52% with 95% confidence. In addition, we find that the slopes of the function f is decreasing in the rank—it is higher for ranks that are more visible. That is, the visibility effect in convex, the more popular an app, the more it gains from it.

The estimators under both W^2 and $\widetilde{W}^4$ are similar for the top 4 bins. The estimators under W^2 show a stronger visibility effect. This is to be expected, as $\widetilde{W}^4$ places more restrictions on swaps, and thus are likely to better isolate exogenous demand shifters.

3.4 Discussion and model refinements

In this subsection we discuss a number of considerations regarding the model and offer refinements to address them.

Competition. Rank shifters resulting from variation in the demand of a different app represent a change in the competitive environment that, in turn, may affect our estimators. Essentially, our

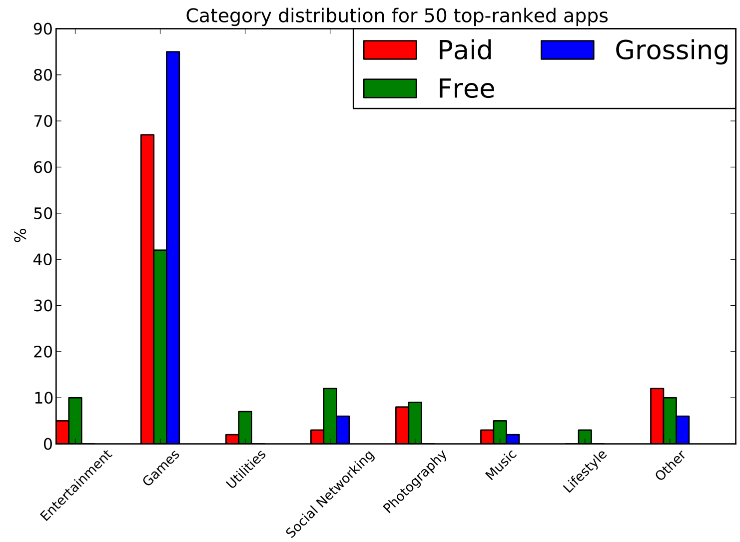

Figure 7: caption

estimators may include substitution effects between apps. This is mostly a concern for apps in the Games category that represent about 67% of the top 50 positions in the top paid chart (see Figure 7). We consider two sources of potential bias.

First, a rank variation may result from a change in the competitor set, when a new direct competitor is released. We control for this with the variable *first appearance* that records whether the top 25 position has seen a new release in the past 24 hours (this time period was chosen after experimentation, accounting for the fast dynamics of the app store). We can further consider first appearances of apps from the same and from different categories. Somewhat surprisingly, we find that the variable first appearance is not significant in any specification, and the estimators are indistinguishable.

Second, the demand for an incumbent competitor may shift in a way that unobserved to us (e.g., a marketing campaign). If this results in a rank change for a no-swapping competitor, we may wrongly attribute the decrease in grossing ranks to a visibility effect, while in fact it should be attributed to competition. Hence, this bias, if exists, should be higher for apps in more competitive categories, such as games (see Figure 7 for the distribution of top-ranked apps across categories). To test the potential effect, we introduce category dummies in our estimation procedure. None of the category dummies is significant in the 10% threshold, out of 18 categories represented in W^2 and $\widetilde{W}^4$. Moreover, the estimators under both specifications are indistinguishable.

Both modifications explored fail to find that are estimation procedure is biased by a competition effect.

Effect of rank change of grossing chart. When one app changes its position in a top-ranked chart, there must be at least one other app changing position. Thus, an app's rank variation in *both* charts, downloads and grossing, may be explained by the rank variation of different app, limiting the potential to demonstrate an effect of rank position on demand.

To circumvent this, we use the fact that top-grossing charts include both free and paid apps, and consider variations in the grossing rank only with respect to apps from the opposite group, free or paid, i.e., for the purpose of this estimation we take the grossing rank of a paid app to be the number of higher grossing-ranked free apps (plus 1), and that of a free app to be the number of higher grossing-ranked paid apps (plus 1). It is assumed here that variation in the ranks of apps from the opposite pool will not affect the demand for the app under consideration, since these do not compete in the same dowload chart. That is, to the extent that rank position affects the underlying demand, we assume that it is popularity charts (downloads) that control this effect, and not the grossing charts.

In-app purchases. In recent years top-grossing charts are dominated by free apps monetizing with in-app purchases. Our analysis centers on the paid app (with non-negative price) in the overall top-ranked list. We focus of the overall top-ranked list since it is by far the most visible chart in the marketplaces and since it includes apps from different categories that allows us to control for competition to some extent (the overall chart is still primarily controlled by games). Compared to free apps that make their revenue from in-app purchases, for which little data is available, the revenue for paid apps is composed

We concentrate on paid apps because of the lack of knowledge on consumer behavior facing in-app purchases that are the sole source of sales for free apps in the computation of the grossing position. Most paid apps feature in-app purchases, but the lion's share of sales is the download price, as consumers are reluctant to spend more money on apps for which they already paid, according to industry sources.

4 Structural Model

In this section we study optimal pricing policies in a stylized model of the mobile app market. To capture the indirect effect identified in the previous section, we consider a dynamic model where past sales increase visibility that translate to higher sales in the present. The seller is assumed to maximize its infinite horizon discounted revenues with discount factor $\eta \in (0, 1)$, under a zero cost assumption, that is reasonable in an information good setting. We comment here that app marketplaces offer sellers a standard revenue sharing agreements, for example App Store keeps 30% of the revenue generated. Therefore, explicitly accounting for distribution costs over the platform will not change the optimal policies.

Our formulation assumes discrete time $t \in \mathbb{N} \cup \{0\}$. Downloads at time t are a denoted by d_t and are given by a function $D(d_{t-1}, p_t)$, where d_{t-1} captures the indirect effect and p_t is the price charged at time t. We denote by $D_d := \partial D(d,p)/\partial d$, $D_p = \partial D(d,p)/\partial p$, and similarly for higher derivatives. The indirect effect entails $D_d \geq 0$, and an inverse relationship between price and downloads $D_p < 0$. We further assume the existence of a maximal download potential $\bar{d}$ such that $D(d, 0) \leq \bar{d} < \infty$ for all d.

The seller seeks to maximize discounted revenues given by

$$\max_{p_1, p_2, \ldots} \sum_{t=0}^{\infty} \eta^t p_t d_t \tag{9}$$

$$\text{subject to: } d_t = D(d_{t-1}, p_t) \text{ for all } t \tag{10}$$

for some predefined d_{-1}.

Interestingly, we find that (9) can be reformulated to match a classic optimization problem studied in the literature of business cycles and chaos as in Benhabib and Nishimura (1985). This class of problems considers an economy balancing consumption and an investment of a single good. If the economy restricts consumption it gain higher levels of the good in the next period that can be used for future consumption, but also to increase future production. Similarly, a developer who keeps the price low can charge a higher price in the future, but also sees an upward shift in her demand function.

The main results of this literature evolve around finding conditions under which the optimal consumption policy in cyclic that parallels with a cyclic pricing policy. We refer the reader to the results in Benhabib and Nishimura (1985) and Baumol and Benhabib (1989) for the mathematical results.

In practice we find that the strategy of pricing cycles is followed by a number of top-ranked apps, mostly ones from big developer companies such as Electronic Arts (EA) and Zynga. Figure 8 shows a time series of ranks for The Sims 3 that is developed by EA and one of the top-ranked apps of all times. During the period depicted The Sims 3 entered 4-5 price cycles, about 1 per month, where the app's price lowered from the standard price of \$6.99 to either \$2.99 or \$.99. The result of the price drops was higher rank that could have contributed to downloads via the indirect effect.

Building on our estimation results, we know that a price cycle policy is effective only if the rank position under the low price is between 1 and 20, which is the case when The Sims 3 sells for \$.99. Since users' price sensitivity varies across apps, we cannot make unifying pricing recommendations.

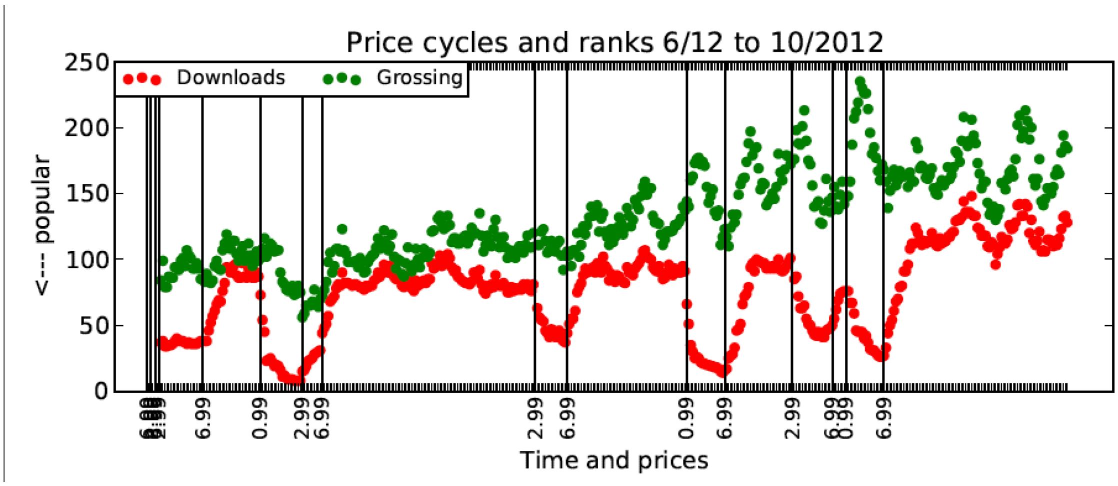

Figure 8: Time series of overall download and grossing ranks for The Sims 3 from June to October 2012. Price changes are marked with vertical black lines, with the price at the bottom.

However, app developers can conduct simple price experimentation to infer their users willingness-to-pay, and make use of our analysis to determine their optimal pricing policy that, as we have shown above, my be non stationary.

5 Conclusions

This paper is among the first to study the economics of the market for mobile apps. We develop a unique empirical model to measure the indirect or visibility effect, while circumventing the endogeneity of explanatory variables. Our estimates show a strong visibility effect for the top 20 apps in the overall top-paid chart. Building on these estimates we find that price cycles may be the optimal pricing policy and demonstrate that this pricing policy is used in practice.

References

Baumol, W. J. and Benhabib, J. (1989) Chaos: Significance, Mechanism, and Economic Applications. *Journal of Economic Perspectives* **3**, 77–105.

Benhabib, J. and Nishimura, K. (1985) Competitive equilibrium cycles. *Journal of Economic Theory* **35**, 284 – 306.

Brynjolfsson, E., Smith, M. D., and Hu, Y. (2003) Consumer surplus in the digital economy: Estimating the value of increased product variety at online booksellers. *Management Science* **49**, 2003.

Carare, O. (2012) The Impact of Bestseller Rank on Demand: Evidence from the App Market. *International Economic Review* **53**, 717–742.

Chevalier, J. A. and Mayzlin, D. (2006) The Effect of Word of Mouth on Sales: Online Book Reviews. *Journal of Marketing Research* **44**, 345–354.

Distimo (2013) *2012 In Year Review.* URL www.distimo.com/full-year-report-2012?

faberNovel (2010) *How to successfully market your iPhone application?* URL http://www.slideshare.net/misteroo/how-to-market-your-app.

Garg, R. and Telang, R. (2012) Estimating App Demand from Publicly Available Data. Mimeo.

Sorensen, A. T. (2007) Bestseller lists and product variety. *The Journal of Industrial Economics* **55**, 715–738.

Surikate and GfK (2012) An insight into iPhone user behaviour within the App Store. URL http://www.surikate.com/en/etudes.html.

Wakabayashi, D. (2013) The Point-and-Shoot Camera Faces Its Existential Moment. *The Wall Street Journal* .

A Appendix: Estimation Results

In the following estimation results * indicates 5% significance, ** 1% significance, and *** .1% significance.

	Estimate	Std. Error	t value	Pr(>|t|)
(Intercept)***	0.0141	0.0017	8.18	0.0000
$\beta^{***}_{[1,5]}$	0.0543	0.0142	3.81	0.0001
$\beta^{***}_{[6,10]}$	0.0307	0.0097	3.16	0.0016
$\beta^{***}_{[11,15]}$	0.0248	0.0067	3.73	0.0002
$\beta^{*}_{[16,20]}$	0.0226	0.0097	2.34	0.0192

Table 1: Estimation results of (8) on $\widetilde{W}^4$ with robust (White) standard errors. The multiple R^2 is 0.003222 (adjusted is 0.002505). The F-test's p-value is 3.683×10^{-9}. The coefficients for bins greater than $[16, 20]$ are not significant in 5% level.

B Appendix: ranking algorithm

We develop a simple estimation procedure to determine who much of the variation in rank positions can be explained by downloads, as well as the weight given to the downloads made in

	Estimate	Std. Error	t value	Pr(>|t|)
(Intercept)***	0.0130	0.0015	8.77	0.0000
$\beta^{***}_{[1,5]}$	0.0526	0.0139	3.79	0.0002
$\beta^{***}_{[6,10]}$	0.0324	0.0079	4.09	0.0000
$\beta^{***}_{[11,15]}$	0.0249	0.0050	5.03	0.0000
$\beta^{***}_{[16,20]}$	0.0265	0.0044	5.96	0.0000
$\beta^{***}_{[21,25]}$	0.0194	0.0044	4.45	0.0000
$\beta^{***}_{[26,30]}$	0.0205	0.0059	3.48	0.0005
$\beta^{***}_{[31,35]}$	0.0171	0.0044	3.91	0.0001
$\beta^{**}_{[36,40]}$	0.0159	0.0059	2.71	0.0068
$\beta_{[41,45]}$	0.0054	0.0044	1.24	0.2163
$\beta^{*}_{[46,50]}$	0.0239	0.0105	2.27	0.0230

Table 2: Estimation results of (8) on W^2 with robust (White) standard errors. The multiple R^2 is 0.00869 (adjusted is 0.007203). The F-test's p-value is 2.2×10^{-16}. The estimators are similar to those obtained for observations $\widetilde{W}^4$, but here more coefficients for higher ranks are significant.

each of the past k days. Following Chevalier and Mayzlin (2006), Carare (2012) and Garg and Telang (2012), we assume that the download distribution underlying the top-ranked charts is a power law distribution. Namely,

$$r^l_{it} = a^l (h^l(D_{it}, S_{it}))^{\theta^l} \tag{11}$$

$$\log(r^l_{it}) = \tilde{a}^l + \theta^l \log(h^l(D_{it}, S_{it})) \tag{12}$$

for different functions h^l, where $\tilde{x} := log(x)$.

We consider the downloads-ranking data of two apps: App C is in the *Catalog* category for which we have daily data for a period of 126 days in early 2013; and App F in the *Food and Drinks* category with daily data of 252 days starting at the end of 2012. Both apps initially sold for \$.99 and app C raised it is price to \$1.99 15 days after the starting date of the dataset. Table 3 reports the mean and standard deviations. For the purpose of this investigation we consider popularity (paid) ranking only.

App	n	Mean daily rank (standard deviation)	Mean daily downloads (standard deviation)
C	126	18.0 (6.1)	15.0 (4.6)
F	253	5.6 (3.0)	81.3 (33.4)

Table 3: Ranks and downloads statistics for apps C and F.

Recall that variation in rank position is affected by multiple factors, both relating the app itself and to factors of competing apps in the same category. Our first estimation model seeks explain how much of the variations in rank can be explained by variations in downloads. Since we do not observe function h^l, we first assume that it is given the number of downloads on the same day, i.e, $h^l(D_it, S_it) = d_{it}$. This is a conservative approach that does not take into account the entire download history.

The regression equation follows (12) and is given by

$$\log(r^l_{it}) = \tilde{a}^l + \theta^l \log(d_{it}) + \epsilon_{it}, \tag{13}$$

where $\{\epsilon_{it}\}$ is assumed to be a sequence of disturbances such that ϵ_{it} and d_{it} are uncorrelated[6]. Table 4 reports the old estimators as well as the adjusted R^2 for (13). We find that variations in downloads explain 30-58% of the variations in ranking positions, as given by the adjusted R^2. Given the variety of factors influencing the top-ranked position, we find this figure high enough to conclude that downloads are key inputs to the ranking algorithm.

	App C	App F
	estimate (standard deviation)	estimate (standard deviation)
$\tilde{a}^l$	4.56 (.22)	5.69 (.222)
θ^l	-.61 (.083)	-.95 (.05)
$\bar{R}^2$	.3	.58

Table 4: Estimators and adjusted R^2 for estimation equation (13). All estimators are significant at the .1% level.

Having established that downloads are key inputs to the ranking algorithm, we proceed to establish better understanding of the function h^l. Namely, we will determine whether lagged downloads play a role in the ranking positions, and if so what are their relative weights, by studying the equation

$$h^l(D_{it}, S_{it}) = w_0 d_{it} + w_1 d_{it-1} + \ldots + w_\tau d_{it-\tau},$$

where τ is to be determined.

To simplify the estimation, we linearize the right-hand-side of (11) using the estimates for a^l

[6] If a visibility effect exists, then this assumption might fail. However, both apps are top ranked on low visibility categories. Our analysis in Section XXX shows that this is of little concern.

and θ^l obtained above

$$(r^l_{it}/a^l)^{1/\theta^l} = w_0 d_{it} + w_1 d_{it-1} + \ldots + w_\tau d_{it-\tau} + \epsilon_{it}. \tag{14}$$

We estimate (14) for both app C and F for $\tau = 1, 2, 3$. For app C we find that none of the lagged downloads has significant effect on the ranks. That is, we cannot reject the null hypothesis that the weight given to lagged moments is zero (the lower p-value across all estimates is .29).

For app F we find that $\tau = 1$; the estimators for w_2 and w_3 are not significantly different than zero. Table XXX reports the estimator for app C with $\tau = 1$. Even in this case, we find that the weight for the current day's downloads is about 3.5 higher than the weight of the lagged downloads.

	Estimate	Std. Error	t value	Pr(>|t|)
downloads0	0.8585	0.0896	9.59	0.0000
downloads1	0.2431	0.0889	2.73	0.0067

Table 5: Regression equation (14) for app F and $\tau = 1$.

To compete the analysis we report below the correlation matrixes between lagged downloads to verify that our results are not biased by strong multicollinearity. Table 6 shows that this is of little concern. Thus, we conclude that same period downloads are key the inputs to the ranking algorithm and this will be used in developing our estimation model in Section XXX.

	App C			App F		
	d_{it-1}	d_{it-2}	d_{it-3}	d_{it-1}	d_{it-2}	d_{it-3}
d_{it}	0.08	0.16	0.17	0.69	0.51	0.35

Table 6: Correlation between lagged downloads, where t is measured in days.

Additional Reading:

Statement of the Publish On Demand Research Theme
http://mediax.stanford.edu/POD/concept

For more information:

- membership
- research themes
- events (conferences, seminars, workshops etc.)

Please visit our website - *http://mediax.stanford.edu*

Like us on Facebook -
https://www.facebook.com/MediaXatStanford
Follow us on Twitter -
https://twitter.com/mediaXStanford
Join us on LinkedIn -
http://www.linkedin.com (search for MediaX at Stanford)
Watch us on YouTube -
http://www.youtube.com/user/mediaxstanford

or contact:

Martha Russell, Executive Director - marthar@stanford.edu

Jason Wilmot, Communications Manager - jwilmot@stanford.edu

Adelaide Dawes, Program Manager - adelaide@stanford.edu

Acknowledgements

Many thanks to mediaX Members and Partners, whose contributions have supported this Research.

Special thanks to the professors and researchers who provided these research updates. Thanks also to Addy Dawes, Program Manager, and Jason Wilmot, mediaX Communications, for editing and putting this report into the format you are reading.