Economic Value Creation in Mobile Applications

Timothy F. Bresnahan, Jason P. Davis, Pai-Ling Yin[1]

ABSTRACT

New mobile development platforms have created an enormous opportunity for applications innovation by lowering the costs of R&D and attracting the largest body of potential app demanders ever assembled. This has created an explosion of app suppliers, numerous and diverse. In this early stage of discovering how this new industry will create economic value, such a wide variety of value-creation experiments is potentially very valuable, a wide variety of market institutions have been created to support those experiments. However, as we document in this paper, the sheer volume of app product entry has created problems for marketing and commercialization, most importantly the challenges of matching consumers to products. The high cost of product matching particularly impacts entrepreneurial app developers. A number of institutions and firm strategies have emerged in response to the matching problem; while these ameliorate the problem at the firm level, many of them worsen it at the market level. We conclude by considering how this situation has impacted the industry's task of discovering economic value and choosing among different app and platform features to make its ultimate contribution to economic growth, and ask how this pattern of technical success and commercialization struggles is different from 20th century predecessors.

---

[1] Stanford University and MIT Sloan School of Management. This research project is based on data collection and analysis over a wide range of data sources. We are very grateful to a number of research assistants who have worked on those datasets, gathered industry information, and joined us in industry interviews. These include Markus Baldauf, Sean Batir, Robert Burns, Jane Chen, Sherry Fu, Osama El-Gabalawy, Carlos Garay, Jorge Guzman, Alireza Forouzan Ebrahimi, Tim Jaconette, Nayaranta Jain, Julia Kho, Sigtryggur Kjarttansson, Xing Li, Derek Lief, Sean Mandell, Laura Miron, Jaron Moore, Yulia Muzyrya, Abhishek Nagaraj, Joe Orsini, Hatim Rahman, Sam Seyfollahi, Melissa Sussman-Martinez, Masoud Tavazoei, Sylvan Tsai, Julis Vazquez, Joon Yoo, and Parker Zhao. We are also very grateful to the many industry participants who have shared their time and expertise with us. Pai-Ling Yin and Jason Davis benefitted from the Karl Chang (1965) Innovation Fund and Edward B. Roberts (1957) Fund.

1. INTRODUCTION

No discussion of the great changes in the innovation processes of the 21st century would be complete without an examination of one of the newest growth poles: mobile applications. The Internet has diminished the role of physical barriers to digital innovation. The rise of electronic platforms and markets has generated huge opportunities for entrepreneurs by lowering the costs of development and distribution of mobile applications. However, growing pains have accompanied these changes in the innovation environment: the explosion in entrepreneurship has created an explosion in competition and an overwhelming choice set for the consumer. Application developers employ a wide variety of commercialization and monetization strategies as they struggle to attract the consumer for whom they create the most value and translate that demand into sustainable revenues. The current matching problem between developers and their customers slows the rate and direction of innovation in mobile application. We characterize the strategies, emerging institutions, and implications for platform competition that result from this matching problem. Our empirical investigation of the supply of mobile apps sheds light on the following more general economic topics:

1. Platform innovation: How is complementary innovation coordinated to create a new industry?
2. Industry evolution: How do markets (and other institutions) form to support exchange in a new industry?
3. Value creation: What will the most dynamic new industry of today bring to economic growth through scientific and technical opportunity?

We choose to focus our study on the app developers and examine the rest of the industry participants from their perspective. A robust literature exploring the technological history of mobile communication devices, platforms, and users, but the app developers – the parties responsible for private and social value creation – are largely unstudied, even in the for-profit domain. They are more interesting for the economic questions than for the commercial ones. The occurrence of the mobile apps platform initiative at a time of generally available data, the focus of many participants in the ecosystem on "big data," and the opportunity to engage in our study before the value creation process has settled down makes this an ideal locus to study an important general phenomenon.

In this paper, we present the theory behind platform innovation and then document the current industry frictions that limit the speed and direction of innovation in mobile applications. We examine the current ranking mechanisms employed by Google Play and iTunes to help consumers find apps. We show that the distribution of app downloads relative to use of apps is highly skewed, suggesting a mismatch between the two. We also show that there is a high degree of churn in these rankings, and we argue that this is symptomatic of developer efforts to manipulate the rankings and/or a reflection of heterogeneous preferences for apps, both of which point to the weakness of the collaborative filtering nature of the rankings. We then document the commercialization strategies of apps and find evidence that the matching problem drives a huge portion of advertising activity. Finally, we examine the multi-homing behavior of developers across Google Play and iTunes. We argue that the low rates of multi-homing and the predominance of well-established non-app firms among the multi-homers suggest that only firms who already possess assets which permit them to overcome the matching problem can port to a second platform. This limitation slows the diffusion of mobile app innovation to consumers and further slows the market tip to a single platform.

The broad economic picture of value creation that we are finding in the mobile ecosystem is reminiscent of earlier information and communications technology (ICT) platform industries. Raw technical progress (faster and smaller computers like mobile devices, faster communications, etc.) has a higher rate of change than applications innovation to create economic value. This has been noted in corporate computing, personal computing, business data communications, and the commercial internet.[2] Like mobile, these earlier platforms had rapid invention in purely technical components and successful exploitation of social scale economies, but slower, though very valuable, innovation in application. The slowness of that commercialization and value creation process in ICT platforms has been one of the leading determinants of the aggregate growth rate of the rich economies in recent decades.[3] Uncertainty about the value proposition for a new technology leads to exploration – not only at the scientific stage, but at the commercialization stage (cite Klepper, etc.) This is typical of general purpose technologies (GPT), where the industry does not necessarily know all the uses or even the main uses at the beginning. Although we document the problems of the industry during this early period, we are optimistic about the eventual resolution of the matching problem in the future.

2. INNOVATION IN PLATFORM-BASED INDUSTRIES

A new platform-based industry always has elements of a General Purpose Technology (GPT). Increasing returns to scale arise at an industry-wide level because some of the components used in different applications are common. By the same token, issues of coordination among inventors of applications and/or of the general components arise.[4] A new platform-based industry can sometimes recombine general-purpose components that are already in existence. This, too, creates economies by avoiding re-invention of the existing components and creates coordination problems by bringing existing suppliers into the coordination loop. However the general components are supplied, a new platform-based industry presents applications developers with a partial solution. It is typically up to applications developers to discover or invent valuable uses of the general technology, found the markets needed for valuable uses, and engage in other social value-creation activities.

In our discussion of the platforms for mobile app development, we begin with the most successful areas so far, the recombination of existing technologies and the creation of new general purpose technologies that dramatically lower the costs of inventing a new mobile application.

The first stage of innovation for the mobile app industry began with the recombination of existing technologies to provide an infrastructure for app development and consumer utility. The rapid improvements in the portability and power of mobile devices combined with the rapid improvements in the networking capacity of the mobile telephone and Wi-Fi networks created a huge opportunity for mobile devices to provide much more utility than simply a communication device. The established consumer familiarity with networked computing also paved the way for rapid consumer adoption of mobile applications. Users already understood the concept of accessing

[2] Cites to Bresnahan and Greenstein(xxx), Bar(xxx)
[3] Cites to Sichel and xxx (xxx).
[4] GPT cites. Contrast to platform pricing cites. XXX

remote sites and downloading software.[5] The increased bandwidth further augmented the ability of the smartphone to become a powerful user interface by allowing much of the storage and processing power to be relegated to the cloud. This allowed the mobile device to become both more portable and more powerful at the same time. The recombination of these existing technologies and the improving savviness of consumers permitted the invention of new GPT components, in particular, the iPhone, iOS, and (expanded) iTunes store, with parallels on Android side.

The impact of all of this recombination and new invention was threefold. First, an application developer can create a new mobile app spending only a tiny fraction of the overall R&D cost of providing it. Much of the R&D cost, including invention of mobile devices, invention of mobile telephony transmission, invention of the commercialized internet, of the cloud, etc., has already been born and is spread over thousands of applications. So, too, has much of the investment cost in infrastructure, such as putting mobile phone cells in place. None of these common R&D or infrastructure investment costs are marginal to a particular app. For mobile applications developers, the fixed costs of offering a working system to users have been dramatically lowered. Second, to the extent applications developers took up this challenge, the economic return to new invention and investment in the platforms and in the pre-existing complementary technologies which were recombined would be raised. Third, the opportunities for applications developers were left largely up to them to discover. We shall return to this theme below. For now, we continue with the positive side of positive feedback.

Network effects imply a second external economy, the "indirect network effects" of attracting both a significant body of demanders and a significant supply of applications so that the platform gets over the hump into viability.[6] In the case of mobile apps, this was achieved by the supply of a modest number of influential apps, which, taken together, were sufficient to attract a significant body of demanders.[7] These were a media store, especially for music, *iTunes,* an app for accessing the web, i.e. a *browser*, an informative *map*, and, after a brief interval, some *Games*. Together with some ease-of-use improvements over existing phones, and some economies of purse and pocket space (phone and music in one device), the supply of these "killer apps" led to an initial expansion in user demand for smartphones that could run apps. This created an enormous market for new apps. The combined impact of the existing mobile telephone system, which created the platform products themselves, and the killer apps generated huge growth in mobile device demand, to date largely smartphones. Today, more new smartphones (6^{th} year of diffusion) are sold worldwide than PCs (28^{th}). For some kinds of consumers in some economies – for example, 20-somethings in South Korea – the diffusion of smartphones has gone farther than the diffusion of television. App developers have access to a very large body of demanders. Advertisers who seek to run ads in apps do so as well.

The first thing that is notable in this story of obtaining positive feedback around the two successful mobile platforms is that it was achieved with remarkably little coordination. There was no widespread contract with developers or users – the mass of developers and users were simply offered an arms-length opportunity.

---

[5] Though there were technical antecedents to mobile applications (games and ringtones on Symbian phones, email and messaging on Blackberries, online stores run by carriers, etc.), none of them launched the mobile app industry due to the lack of the necessary critical mass in use and developers necessary to generate positive feedback loops.

[6] n/w fx

[7] This is not the only way to get over the hump into viability. A large demander, such as a Defense Department, can attract sufficient suppliers to start the positive feedback loop, or a platform sponsor can coordinate the joint attraction of many demanders and suppliers. As with many other commercial computing and communications platforms, however, the leap of mobile platforms over the hump was achieved by the attractiveness of a few "killer apps," i.e., apps attractive enough to give users a motivation to buy the product.

The rapid emergence of that many demanders, together with the very low barriers to entry created by the platform providers, has led to a rapid and very substantial expansion in the number of overall apps. Figure 1 shows the dramatic growth in both iOS and Android apps. As is easy to see, the iOS growth starts earlier, as the iPhone was widely marketed before any Android phones.

Figure 1: Apps offered by platform over time

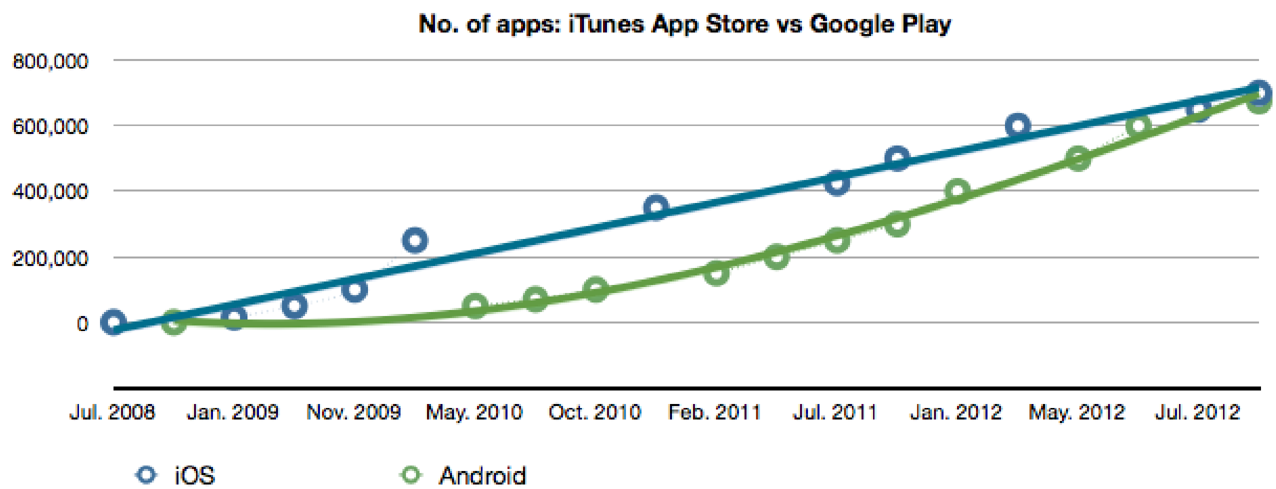

Source: InsideMobileApps.com[8]

Android apps, measured by number, have caught up to iOS apps. One must be careful, however, drawing any economic inference from a count of apps; the marginal app, and probably the majority of apps in the raw count are marginal apps, has not been downloaded or used by customers. We will revisit the question of the size of the supply of Android vs. iOS apps below.

This industry is still in its early stages. The big driver for huge growth in this industry was the establishment of a platform that drastically lowered the costs of entry into mobile application development: it provided both the R&D and the distribution system (iTunes/Google Play stores) at much lower cost than would typically be faced by an entrepreneur. That has, as we just saw, enabled an explosion of entry by app developers, a strong reflection of how low the incremental cost of an app has fallen.

The pace at which the industry creates social value depends not only on those low technical costs but on the rate at which new, high value markets are founded and at the pace of innovation and competition driving up product quality and driving down prices in new products enabled by the technical opportunity in the industry.

3. MATCHING APPS TO CUSTOMERS: APP STORE RANKINGS

Traditionally in new technology industries, early stage pre-dominant design is determined by nerd to nerd marketing with technically fluent customers. Today, in this mass consumer market, the decision makers for the rate and direction of innovation are not technically savvy. Moreover,

[8] Accessed July 19, 2013: http://www.insidemobileapps.com/2012/09/26/tracking-growth-the-itunes-app-store-vs-google-play/

the extremely rapid growth of user base means that distribution for demand for apps can be volatile. Commercialization for entrepreneur is heavily defined by this dynamic. The entrepreneur can benefit from rapid adoption. However, the ability to capture the "right" consumers in that dynamic is unclear, leading to lots of expensive investment in capturing all consumers.

The same platforms that lowered the cost of entry are also the source of the largest costs to developers: the platforms have not been able to provide adequate institutions to help match consumers to the overwhelming product offerings. There current mechanism for solving the matching problem are various "top ranked" lists, which are forms of collaborative filters that aggregate download or revenue information on the apps to determine the "top" apps, with a number of "editor picks" lists thrown in. Our empirical analysis will show that these rankings fail to reflect the value created by the app and match heterogeneous consumers to apps. As a result, a host of interesting services have arisen to try and solve commercialization and monetization problems for apps struggling to get noticed in the clutter of apps.

3.1 iTunes Store

Apple tightly bundles the services of the iTunes Store as a distribution and app-discovery market with the technical products that make up its platform. This is consistent with Apple's overall plan of providing a controlled "stack" of components and services, and to vertical integrate into all of the general-purpose components.

From one perspective, the iTunes store is a roaring success. The applications store was added to iTunes with 500 apps in July, 2008; by May 2013 users had downloaded 50 billion apps.[9] Apple's tight control over distribution gives it control over developer revenues, and Apple charges distribution fees of 30% of revenues. Through May 2013, Apple had paid developers over $8 billion.

From another perspective, the iTunes store is much more problematic for developers. The mechanisms by which users and apps are matched are limited. A user can search for apps by keyword. A user can look at "top lists," such as "top free apps" (or top paid or top-grossing, i.e., highest revenue.) A user can also look at "top lists" within broad applications categories. Finally, a user can arrive at the iTunes store knowing the app they want, either by following a link from another website or by remembering the app's name. The "top lists" are, from a user perspective, a collaborative filter. As each user arrives at the site, the top lists show those apps earlier users have chosen to download. Our understanding is that Apple deals with these difficult tradeoffs by displaying, on the "top" lists, apps which have been most downloaded in the previous 24 hours. This makes the collaborative filter very responsive to current demand conditions. The iTunes store also offers user-opinion-based collaborative filters, with comments and ratings (one to five stars) written by earlier users and read by later ones.

3.2 Google Play Store

Google, the supplier of the Android market, similarly runs an online app store called Google Play.[10] However, comparatively open-systems Google does not bundle services of Google Play with the Android platform, so developers can, if they choose, distribute apps themselves or use a different online store, such as Amazon, which has an Android apps store. Despite the possibility of developers going elsewhere, Google Play has caught up to iTunes' number of app offerings (see Figure 1).

---

[9] See Rowinski (2013).

[10] Before March, 2012, Google maintained separate app, music, and book markets. Since then the former Android Market is merged with Google Music and Google eBookStore and called Google Play.

Google Play has a slightly more complex set of top lists than does the iTunes store, mostly because there are more lists, including Top Paid, Top Free, New Apps and Trending Apps both overall and within categories. Another important difference is that the downloads-based collaborative filter is based on a longer window: our understanding is that "top" apps are the most downloaded over the preceding eight days (vs. one day for the iTunes store.) Otherwise, the structure is basically the same, a mixture of user-action-based collaborative filter (recommending apps recently downloaded for download) and user-opinion based collaborative filter (comments and ratings).[11]

Developers, especially developers who do not have a pre-existing connection to their customers, express considerable frustration with the effectiveness of the app store collaborative filters. Several features of the environment make effective implementation of a collaborative filter difficult. Users typically search from mobile devices, limiting their ability to browse through long lists. New apps are constantly being submitted, so early users would have to choose the best apps quickly if the collaborative filter can find them for later users. Further, most users do not have a large number of apps on their device, so there is limited information to fine-tune the recommendations tailored to a particular individual users. Apple is only channel for distribution, which exacerbates the problem of ranking being the dominant matching mechanism. Google being open and the existence of alternative app stores seems not to be the solution to the hegemony of a single ranking provided by iTunes, just more clutter. It is surprising that Google's expertise in search also does not solve the problem.

The incentive then arises for developers to game the rankings by purchasing downloads to rise in the rankings, which then makes interpretation of downloads as demand suspect. Markets have emerged for inefficient purchase of rankings (Tapjoy, "junk" downloads). Firms have emerged to consult developers with app store optimization strategies. Developers have allocated resources to purchase advertising to promote apps instead of attracting marketing dollars from products outside of the industry. Alternative ranking and rating systems outside the app stores have emerged to better help match consumers to apps, based on niche interests or social networks: Facebook is now well placed to combine social networking with app advertising and app filtering.

The most important implication of this phenomenon, we believe, is the change in the relative position of app developers who already have a marketing connection to their *app* customers, typically because those are already *online* or *product or service market* customers as opposed to new entrants.

We will discuss this behavior more below.

4. DATA

Our tables are built on three main data sources. Each has – and this is an oddity of the particular industry we study – its sampling frame built around <u>products</u>. There are no industry-wide data with the more usual economic sampling frames of firms or markets.

4.1 comScore

---

[11] Apart from the star ratings and reviews, Google Play also has apps recommended by members of a user's Google Plus "circles," i.e. their Google Plus social network. The importance of this improved collaborative filter is limited by the limited penetration of Google Plus into widespread use as a social network.

We utilize the "mobile metrix" dataset from comScore. Like other comScore products, this is based on a panel of users, in this case a panel of approximately 10,000-12,000 US adult users of mobile devices. The bulk of the panel is two subpanels, each approximately 5,000 users, one with Android phones and the other with iPhones.[12]

The underlying fundamental data are about each panelists' possession and usage of apps.[13] However, the data come to us aggregated to the product*platform*month level.

The sample of products (apps) on which data are available meet a minimum usage test for each month. As a result, apps enter and leave the reported sample month to month. For each platform, iOS (iPhone) or Android (smartphone), comScore includes data on the app only if it is used on that platform by more than 5 (at least 6) unique users.[14]

An app is not exactly a product (observation in the data set) but it is close. For some apps, comScore has aggregated distinct versions into a single "property," typically because (1) they view the apps as different version of the same thing, (2) the app supplier views them as the same thing, or (3) the app supplier sells ad space in the different apps as a single ad product. Often, the merged "property" includes both the free part and the paid part of a pair of "freemium" apps, though sometimes free and paid are two separate properties.

It is very difficult to identify the same app on the two different platforms in this industry. In this dataset, comScore staff members (working with developer clients when available) manually identify the same app on iOS and Android. Three interesting issues arise with this. First, this is a continuing project, so not all apps have been processed yet. Second, comScore does not depart from its used-by-enough users standard platform by platform. Thus an app may be *available to users on both platforms* but only included in the comScore data on one platform because it does not have many users on the second. Third, comScore assigns a common, sensible name to each app which is common across platforms. This means that a multiplatform app can have a comScore name which is different from its name on iTunes store and also different from its name in Google Play. An app which is only listed on one platform on comScore can also have a comScore name which is different from its name in that platform's app store.

For each month*platform*(included)app, comScore compiles a number of quanta. These include estimates of the total number of unique users for each of those apps during that month, the total number of minutes for which the app was used by all users that month, the average minutes per visitor (which is the ratio of total minutes to unique users), and the average daily visitors. The average daily visitors is calculated by taking the average, over all days in the month in question, of the number of unique visitors that the app had in a single day. For example, if the universe of mobile app users consisted of two users, one of whom used a particular app every day in April, while the other used it every other day, then the "average daily visitors" measure for that app for April would be 1.5. All of these data are projected to the entire US based on a set of comScore weights for their panel.

We have looked at the joint distribution of all of these different demand metrics and concluded that there is only two dimensional variation in the cross section of apps in a given month. The size metric, unique visitors per month, and the engagement metric, average minutes per visitor, together

---

[12] There is also a smaller subpanel of iPad users, which we do not use in this paper.

[13] There are also data on panelists' visits to mobile websites, but we don't use these data in this paper.

[14] A second criterion could also lead to inclusion, if the app developer has implemented a comScore provided Software Development Kit (SDK) that includes a piece of software in every copy of the app. This piece of software then reports to comScore whether (inter alia) the user has used the app. If than 11,000 unique US users have used the app during the month, it is included. In the months we examine, this second criterion does not appear to lead to the inclusion of any apps that would fail the first test.

can explain almost all of the variation in the rest of the size measures. Both measures are correlated with the two remaining ones, average daily unique visitors and minutes per month, but are very close to uncorrelated with each other (.06). Thus the raw data yield a simple two-factor model of product "size" that we use below.

Sometimes the total unique users measure is converted by comScore into the "reach" of the mobile app. An app's reach is defined as the percentage of all "potential users" of the app who actually used it during the month in question. This is given by the ratio of comScore's estimate of the total number of unique visitors to the app on a particular platform during that month to the total number of users on the platform during that month.

Table 1 shows some monthly statistics evolution of unique apps by platform as reported by comScore between September 2012 and January 2013. Since the comScore panel, census projection methodology, and apps who use the comScore SDK are in flux over time in this very new industry, we are skeptical of reporting any in the data over time. That said, that five month period was one of comparative stability in the definitions and there does not appear to be much change over time in the relative magnitudes across columns of these figures within a month, so hereafter we will focus on the data from January 2013 in our analysis.

The table reports six metrics for each month. "All" is the total number of distinct apps in the data, 1301 in September 2012. "All on Android" (AOA) is the number that are available on the Android platform, and "Exclusive on Android" (EOA) is the number of apps that are available only on the Android platform. Of course, "All on iTunes" (AOI) and "Exclusive on iTunes" (EOI) are defined similarly. Finally, "Multi" is the count of apps in the comScore sample found on both platforms in the month.

Table 1: Number of unique apps by month
AOA=All on Android, EOA=Exclusive on Android,
AOI=All on iTunes, EOI=Exclusive on iTunes

	All	AOA	EOA	AOI	EOI	Multi
2012m09	1301	715	439	862	586	276
2012m10	1243	710	434	809	533	276
2012m11	1202	691	408	794	511	283
2012m12	1203	727	433	770	476	294
2013m01	1231	793	507	724	438	286

In January 2013, users in the comScore panel have used 1,231 distinct apps in total, 793 and 724 of which were available on Android and iOS respectively. Of the Android apps, 507 were exclusive to that platform, while 438 of the iOS apps were exclusive on that platform. Multihoming apps comprised the remaining 286 apps. While these figures are not representative of the absolute sizes of the two platforms they are informative about developer's platform choice.

In Table 2, we report simple descriptive statistics for the size measures. One can immediately can see that all variables are heavily skewed. For instance, on average an app is used 1,993.24 million minutes per month. However, the standard deviation is more than 8 times larger than the mean. The maximum unique monthly visitors for an app is 3-4 orders of magnitude larger than the minimum.

Table 2: Summary statistics of key variables
1517 app-platform observations, January 2013

	mean	sd	min	max
Unique monthly visitors (000)	1289.92	4324.60	11.12	60805.94
Avg. unique daily visitors (000)	300.51	1739.01	0.46	38657.62
Total monthly minutes (millions)	1993.24	16289.40	0.26	391802.25

4.2 App Annie

App Annie is a market research business. They write "App Annie tracks your apps' metrics and has the best app store data to help you make smart business decisions." The App Annie app metrics we use in this paper are download metrics for the online stores (iTunes and Google Play). App Annie copies the rankings of top apps from each of the online store each day. For each day we obtain the ranking of the top 500 free, top 500 paid and top 500 grossing (revenue) apps for iPhone and Android phones. For Android, there is also a 'top new free" and "top new paid' ranking each day.

The two online stores clearly have slightly different ranking algorithms. Neither online store publishes its ranking algorithm, but there is, of course, constant discussion of the algorithms among industry participants. Our understanding is that the top apps on iTunes is based on downloads in the past 24 hours, whereas the top apps on Google Play are based on downloads in the past 8 days.

4.3 App Questionnaire

Finally, we have ourselves undertaken a survey of apps. We created a questionnaire, focused on apps' monetization strategies, categories, content and presence on multiple platforms, and asked research assistants to download apps onto their mobile devices and then fill out the questionnaire. The main point here is that there are a number of critical questions about app developers and apps – especially the monetization strategies, which are public because they are visible in the app to any users but which are not collected anywhere else (they are not observable on any publicly available websites).

We have filled out our questionnaire for each of nearly 5000 free apps (2,281 Android and 2,713 iOS) starting on June 25, 2012. We have two sample inclusion criteria. First, we attempt to survey the most popular apps. This is easy on Android, since Google Play reports, for each app, what total cumulative downloads bin it falls into. We define popular on Android as an app which has had at least 500,000 downloads. It is rather more difficult on the ITunes Store, which has no similar data publically visible. For iTunes, a popular app is one which has been on the AppAnnie rankings list for at least 10 days at any time. We continue to survey popular apps today, but the samples used in the tables in this paper are for apps that were popular anytime since the start of iTunes and Google Play through January 2013.

Our second criterion for inclusion in the questionnaire is if the app is in the comScore data. Surprisingly many apps are in the comScore data, i.e., used by quite a few people but not "popular." This could arise because comScore is a measure of current usage while our definition of "popular" is, at least for Google Play (and intended to be for ITunes) a measure of cumulative downloads. But it is more likely that it reflects the natural sampling uncertainty. In both datasets, we have pushed our definition of "popular" far out into the "long tail" of apps, and comScore has also pushed its definition of "widely used" far out.

For each app, the questionnaire asks whether the app utilizes advertising and what is being advertised, as well as a number of questions about the frequency, duration, and format of the ads. We also ask a number of other questions not used in this paper. Furthermore, we ask the research assistant to look at the developer's website to see if the app is present on the other platform and whether the same developer has a different app on the other platform. Similarly, we ask the research assistant to learn whether the app is part of a "freemium" pair. Our examination of the developer websites is the closest thing to a firm-level dataset on app suppliers in this industry.

The biggest disadvantage arises because we are unsure, on the iOS side, exactly what function of downloads we are using as a cutoff. If ads are targeted, then our data on advertising content is conditional on our respondents' behavior (i.e. students who download many apps per week) and thus not be representative of the market. However, from conversations with advertising firms our understanding is that targeting at the individual level is not commonly used as of today.

5. APP SUCCESS IS HIGHLY CONCENTRATED

Many theoretical models of platform success simply count apps. And the mobile industry tends to follow that, noting that there are approximately ¾ of a million apps available for each of the two largest platforms, with significantly fewer for Windows Mobile or Blackberry (and for the tablet-format of both iOS and Android.) Yet it is not obvious as an empirical proposition that simply counting apps is a good way to think about the contribution that aggregate app supply makes to either the competitive success of a particular platform or to the overall value of all platforms to the economy. Some apps may be much more important in delivering value to consumers and thus much more important in creating (competitive or growth) value for the platform. In this section we examine the cross-section distribution of app demand; in the next section we look at some simple short run dynamics.

To consider the cross section question, we looked in the comScore data to examine the size distribution of app demand. For each of the two larger platforms, the comScore data contain a number of size measures.[15] As we noted in the data section, there is very substantial correlation among a number of these measures, and we are thus drawn to examine only two of them: how many users have the app on their device and use it and how many minutes a typical user has the app open. We begin by looking only at total unique monthly visitors for January 2013. This is shown in Figure 2.

[15] The data we are looking at are for smartphone apps. The ComScore data also have, for a somewhat shorter time series, data on iPad apps. A number of other industry sources attempt to measure more size variables than we have here, notably app revenue.

Figure 2: Size distribution of apps based on total unique monthly visitors, January 2013

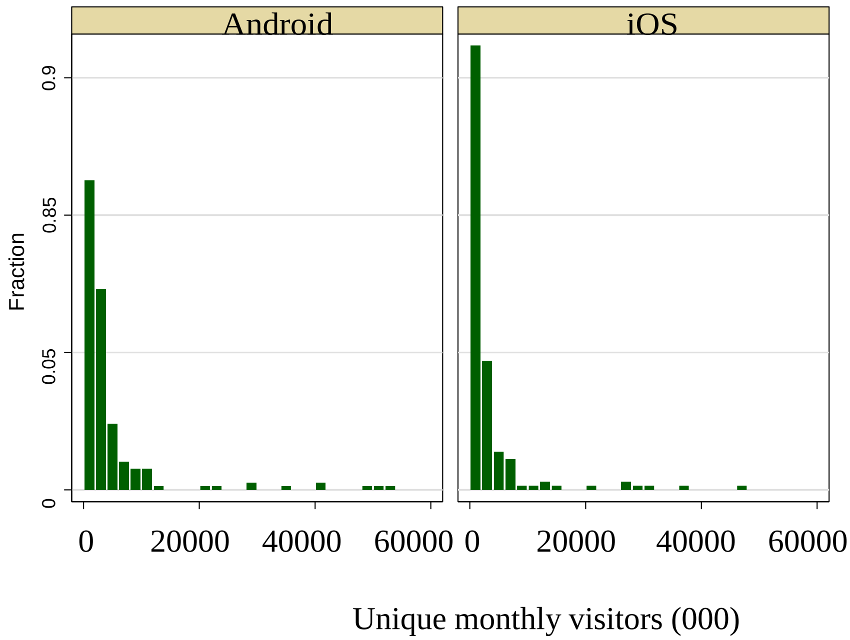

Fixed-width bins, scale for X-axis is thousands of projected US smartphone users.

Both platforms exhibit a high degree of skew in the size distribution of app demand. For each platform, the top 20 apps account for approximately 80% of demand. (Not the top 20%, the top 20).

We can learn a bit more about this considerable heterogeneity in app demand by looking at the joint distribution of the number of users and of the number of minutes of typical use. For this purpose, we look at iOS smartphone apps in the same month, January 2013. We use comScore's reach metric and average minutes per visitor as the two size measures in Figure 3. In this figure, every app ("property") is a dot. The two measures of size are not highly correlated ("reach" is very highly correlated with "unique visitors" and with "total monthly minutes, and none of these is very highly correlated with minutes per visitor in the cross section of apps.)[16] Industry people, not surprisingly, think of this as a distinction between apps which deeply engage the user (y axis) and apps which have mass appeal to a very large number of users. It is clear that these are not the same object. Nonetheless, it is pretty easy to see that there would be a very highly skewed distribution of the attractiveness of apps whether that is measured by reach, minutes, or both.

[16] This feature of the figure would stay the same if we switched to Android or to another month. One important detail would change: this figure is taken from Apple's ill-fated experiment with kicking Google maps off of iOS. Earlier or later time periods and the Android platform would show that app much farther to the right.

Figure 3: Joint distribution of minutes and reach

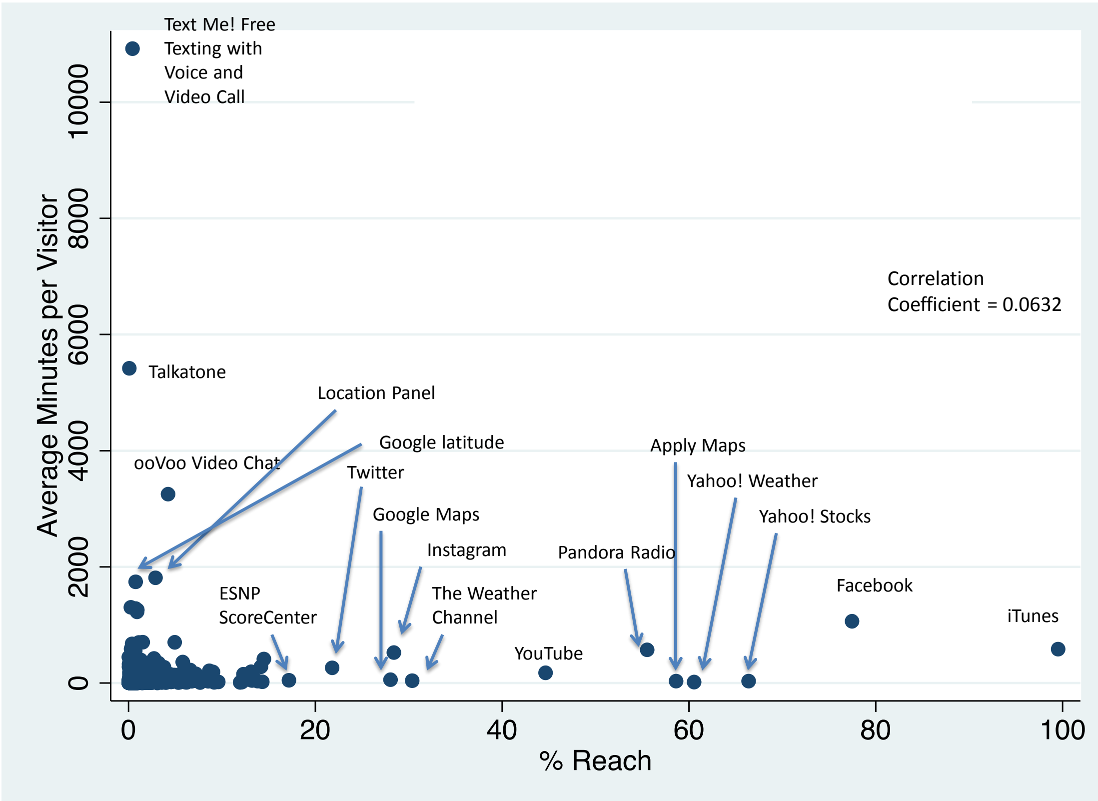

At least one source of variation in the size distribution, pricing, can be descriptively examined. In Figure 4, we look, again for the iOS platform in the same month, at the joint distribution of reach and of price. This is once again based on the comScore data, which we have linked to the iTunes store in order to collect the prices.[17] As a threshold point, one can see from this figure that free apps make up a very large fraction of total app distribution. Paid apps, in turn, if they are high enough up in distribution to be labeled in the figure, are most frequently games.

[17] This matching process leads us to drop a small number of apps for which we simply cannot figure out which comScore app is which ITunes store app. When there is both a free and a paid app in a “freemium” property in comScore, we use the price of the paid app in the figure. When comScore reports the free and paid halves of a “freemium” pair separately, they both show up in the figure.

Figure 4: Joint distribution of price and reach

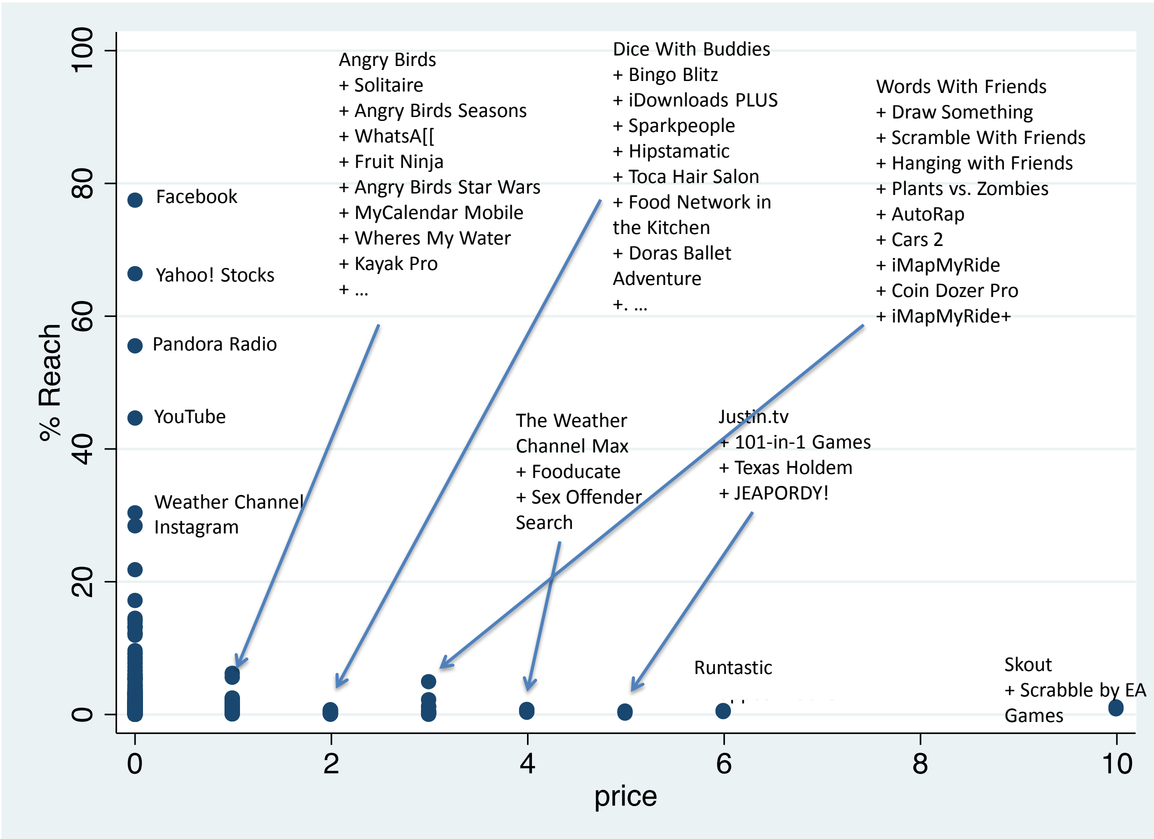

Why the enormous skew in app demand? First, some apps are simply more useful, cheaper, or better programmed and marketed than other apps. Second, among equally attractive apps, some got to market earlier than others – and quite possibly, some very important apps are still growing in our snapshot figures. Finally, the most important distribution and app-discovery mechanisms are the online stores, and those add an element of positive feedback to app demand through the online ranking system. No reader of Sorensen (20xx) will be surprised by this last point: the app stores are a “greatest hits” recommendation system, and the size distribution of apps shows a few greatest hits.

6. SHORT-RUN DYNAMICS

If the reason for the highly skewed distribution of app demand were that there are only a few apps that have been creating most of the value for consumers, we would expect those apps to dominate over an extended period of time. In this section, we present analysis of the churn in app rankings at different time scales and for both Android and iOS apps.

We actually make three changes relative to the last section. (1) We switch from a descriptive snapshot cross section to a descriptive short-run dynamics framework: churn. In this, we follow Waldfogel (2013) and ask the question “how many apps are on the top 10 list.” This is unlike the question of how many top 10 economics departments or business schools there are: at any given moment there are only 10 apps on the top 10 list. However, if we look at two moments, there may have been some turnover. If, for example, (exactly) two apps that were on the top ten list at time s are off it at time t, 12 apps will have been on the top 10 list at the two times. (2) Because we switch from the quantitative demand data of the last section to the rankings of this section, our analysis necessarily becomes more nonparametric. Relatedly, the data we use here are basically administrative records, so the sampling error, representativeness issues, etc. which may arise in the quantitative demand data are less important. Finally (3) we switch from the stock-of-apps-in-use frame of the last section to a flow-of-apps-downloaded in the current one. This is because the administrated records data we use are downloads ranking data gathered from the iTunes store and Google Play by the recommendation site AppAnnie. AppAnnie has been collecting the daily download rankings data over time since the app stores opened. While the iTunes store opened before Google Play, in the figures that follow we restrict attention to the more recent period in which both stores were open. The distinction between stocks and flows may be smaller in this industry (at its present rate of development) than it might be in a more mature industry. The rate of growth of new users and of new replacement devices at existing users is very high, so the stock does not lag all that far behind the flow.

More precisely, using daily rankings from App Annie we determine the volatility of a ranking from two dates s and t as follows. We consider the top 50 apps in both rankings, Rs and Rt. Next, we count the number of apps in the union of the two sets. Table 3 contains the average values of D for several interval lengths (t-s) over a long period of time and for a number of different "top" lists for each platform.[18]

We show statistics for each ranking method for the daily top 50 rankings of apps. Since there are no ties in the data all the reported means all lie between 50 and 100. To read the first number in the first row (Android top paid, t-s=1), for Android paid apps, there are 51.35 apps on the top 50 list over the average two-day period. That is, if we look at Monday and Tuesday, we would expect 1.35 apps, on average, to have fallen off the top 50 list.

What do we see in Table 3? As a broad, general point, there is a good deal of churn and turnover. This is perhaps what you would expect from a very new market. The churn is in general pretty high frequency, with quite a bit of the turnover that happens in a full year (right column of the table) occurring in a month for both iOS and Android. (By construction, all of the Android "new" lists have churn which happens within a month.)

A second point is that after either a day or a week the turnover on each of the Android "new" lists is significantly higher than the overall top paid and top free app lists. This is not all that surprising, either, new apps are experiments, and they mostly fail. Many developers tell us that the Android "top" lists are based on an 8-day history of downloads. Thus, after about a week even a developer who has been trying to "buy" downloads may well know the experiment is not working.

Third, if we look at either the top free or top paid list, there is more turnover of top apps on iOS than on Android. This difference is evident at t-s=1, and continues to grow over time. After a year, for example, just over 50% of the top free Android apps (78.15-50) will have departed, while over 80% of the top free iOS apps (91.24-50) will have departed the top lists. Developers tell us it is more important (and more effective) to "buy" downloads on the iTunes store, because it uses a shorter history of app downloads to form the top list.

Had we looked at a shorter or longer top list – say top 10 apps or top 300 – these findings would have been qualitatively if not quantitatively the same.

[18] The "top paid" and "top free" lists rank downloads. The "top grossing" apps are ranked by revenue. Most top grossing apps have in app payments, and the IAP revenues plus the original payment for the app (if any) are used to create revenues. The "top grossing" list is thus a bit closer to a stock measure. No analog of the two Android "new" lists exists for iOS.

Table 3: Average number of distinct apps in top 50 daily rankings between 01/01/2012 and 06/28/2013

Days in between rankings	1	7	30	90	180	360
Android						
- top Paid	51.35	56.59	61.05	64.56	67.77	72.11
- top Free	51.26	56.00	63.11	68.99	74.60	78.15
- top Grossing	51.56	56.89	64.22	74.32	81.38	91.08
- top New Paid	53.59	69.28	99.69	100.00	100.00	100.00
- top New Free	53.98	72.25	99.91	100.00	100.00	100.00
iOS						
- top Paid	53.64	61.30	69.02	75.29	80.24	83.86
- top Free	55.65	68.18	79.58	85.31	88.24	91.24
- top Grossing	53.96	57.64	62.91	69.18	75.96	84.38

There are, of course, a number of possible explanations for these simple descriptive statistics about churn. New and better apps could be constantly being introduced, displacing old (90 days!) outmoded competitors. Alternatively, new apps, not necessarily competitors of the apps that used to be on the top list, could be most of the flow of downloads while earlier hot apps sit on the stock of phones, having already completed their diffusion. Developers and other market participants have not proposed these stories, and they would need *extremely* rapid dynamic competition or diffusion processes: on the iOS side, 40 to 60 percent of the top fifty at any given date are out of the top 50 a month later.

The churn figures are also consistent with another type of very powerful dynamic competition, competition across all apps, whether substitutes in functionality or not, for consumer attention on the app stores. Developers seeking to "buy" a position on the top lists for their apps by advertising or "incentivized" downloads are competing with one another for a limited resource, the top of the top lists and the attention it brings from potential users.

Heterogeneity of demand could also explain churn in a single index ranking. If preferences are highly variable and widely distributed, the growth of customers over time will cause the actual distribution of preferences to change. In this case, a single collaborative filter ranking that averages over all consumer preferences is still inadequate for matching heterogeneous demand groups to their highest utility app.

7. THE ECONOMIC RETURN TO THE DEVELOPMENT OF NEW APPS

In this section, we take up the factual question of how developers seek to earn an economic return on their development effort. Our primary concern is understanding the formation of new markets, a key step in the creation of value out of a new platform industry. As an incidental payoff, we will be able to address some management-normative questions about the "monetization" of apps.

While mobile apps are a general purpose technology and thus might have a very wide range of uses, three main ideas about the way they might become valuable informed much early app development.

1. Many apps might be part of a cluster of entertainment services consumed on mobile devices. Games played on the mobile device are an obvious example.

2. Many apps might remove life's annoyances from users' lives. Maps are an obvious example.
3. Apps provide a dramatic new advertising medium, with the ability to condition advertising on a user's location as well as on many other targeting data.

As always with new IT platforms, mobile apps were also widely described as having the potential to "disrupt" *everything*, which is, as a practical guide to investment, the same as nothing. As the vagueness of these broad categories suggest, app supply heterogeneity shows a good deal of heterogeneity in the way the developer seeks to earn an economic return if the app is successful.[19] More specifically, developers can choose from a host of different monetization strategies. We begin with static strategies which create an economic return to the app now.

7.1 Charging the user

Some apps, so-called "paid" apps, are bought by the user at the time of download. These apps pick up on the themes of providing valuable entertainment or removing an annoying hassle from the user's lives. For a number of reasons, however, straight up paid apps are however, not the most common revenue-gaining strategy even when the user ultimately pays.

Two different strategies have a try-before-you buy structure. "Freemium" apps are free when the "lite" version is first downloaded, but limited in some way. A "pro" version is available at a positive price, and typically can be bought from within the free app itself or through a link within the app sending the customer back to the app market where the customer can download the paid version.

Similarly, "in app purchasing" lets users buy complements to the app as they are using it; in a warlike game, IAP could be used to buy a bigger sword or more lives for the player's avatar, for example. In app purchasing and freemium models are not mutually exclusive. Many "pro" versions also entice in-app purchases. The widest use of IAP is in games – if the game is diverting, or even addicting, the user pays.[20]

The use of try-before-you-buy strategies in a new product area is not surprising. It is even less surprising when you consider the weak institutions of the mobile platforms for matching buyer and seller; buyers can be much more easily enticed if they do not have to pay upfront.

Other apps use a subscription model, a recurring monthly or yearly in-app purchase that could be managed by the developer's payment system or the app stores itself. It is common for media related firms to offer apps such as Pandora, Hulu, and Netflix on a subscription model that include complete passes to unlimited media consumption. Matchmaking, personal finance, anti-virus and even navigation apps have also called for users to subscribe. App developers may be able to skirt the rules of transferring a portion of their app revenue to the app marketplace if they ask for the payment of the subscription dues outside the app and use a pre-established username/password for authentication. The carriers have even sponsored some subscription apps that allow users to pay

[19] We have also met a substantial subset of developers with no interest in profit whatsoever. There are a number of tech savvy developers who build apps as "art" or tools to make their own lives easier. These developers may be anti-capitalist or simply too well paid by their day job developing software to care about the app profits. They are primarily motivated by the validation of their users using the app they built, the experience of building an app, or the simple utility (including oddity) of the app itself.

[20] It is not uncommon for apps that monetize through in-app purchases to have dual currency types, in an effort to spur greater levels of monetization. For example, one form of currency may be easy to obtain through continued gameplay while another form of currency may be more difficult. Paying the developer through an in-app purchase makes acquiring the scare currency much easier. Further investigation could help us better understand the in-game currency dynamics, especially as they relate to dual-currency games.

through mobile phone bills. However, this practice is much less prevalent in the United States than in other countries including Japan that generally have higher rates of app monetization.

Once we have discussed the other main forms of app monetization, we will present some simple statistics on the prevalence of all the forms.

7.2 Advertising Supported Apps' Monetization

There are also a number of ways that apps can be advertising supported. Both Google and Apple offer "ad brokerage" services to app developers. The developer creates a space in which an advertisement can run – before the app loads, across the top of the screen, between stages of use, etc., etc., etc. – and the broker sells that space to potential advertisers. The app developer can be paid either like traditional media firms, i.e. by the number of users who are shown the ad, or like traditional online firms, i.e., by the number of users who immediately take action in response to seeing the ad. In addition to the in-house advertising brokerage services, there are a number of independent ad brokerage networks.

In addition to the wide variety of organizational structures for advertising, there are also a wide variety of ways to display advertising. A partial list would include text ads, banner ads, "click-to" ads, expanding banner ads, video ads, rotating banner ads, and interstitial ads, i.e. ads which are displayed full-screen during a pause in use of the app, such as between rounds of a game. Interstitial ads are like TV or radio ads, taking over the user interface for a brief time, while the other forms of ads are like newspaper or magazine ads, appearing over or under or beside or in the middle off the app. As innovation by ad-tech companies continues, regular interstitials evolved to offer brands an immersive near virtual reality experience for consumers to explore new products. Additionally, simple banners evolved into pop-up push notifications appearing in the status bar of a phone, which some in the industry have termed "malware". Ad options vary by particular platform. Google allows push notification advertising, while Apple's tight control over the mobile phone software development process leads them to shun this advertising practice.

Additionally, as mobile phone screens are limited in size, erroneous clicks become more prevalent. This gives developers the opportunity to undertake less-than-consumer-friendly strategies; some have placed ads near the portion of the screen where the user might click to play a game. This leads game players to erroneously click on ads, boosting the click through rate and making the ad more lucrative. Other developers go down a very different path, allowing one full use of the app before showing ads – an indirect version of try-before-you-buy.

Finally, ads may be more or less targeted to the user. In principal, targeting could be based on a large number of different actions on attributes of the user, including their location, their past purchases, the content of their communications, etc., etc. Some observers believe that targeting will lead to a dramatic reduction in trade frictions, as users are presented with valuable product offers at just the right time, place, etc. to create gains from trade. Other observers believe that targeting is a large scale loss of privacy for consumers and represents an important loss of consumer sovereignty. Both sides are largely forecasting the future, not describing the present. The scale, scope, and form of targeted mobile advertising are changing; targeting today is typically quite primitive. While some advertising is integrated with the users' past behavior, location, etc., many other ads are much closer to broadcast than targeted. A number of complementary inventions are going to be needed before there is much effective targeting in mobile advertising. App developers are going to have to define and create audiences, the way traditional media do today.[21] Ad networks, including those which are part of the large platform firms, are going to need to work out what to measure about users and how to sell that knowledge, and are going to have to

[21] A number of a brokers and advertisers have told us that making the target audience (in the sense an advertiser cares about to whom they send messages) part of app design lies largely in the future.

figure out how to divide the rents associated with targeted advertising with developers. Users are going to need to decide what to react to, after learning (this, too, is in the future) what kinds of offers they get from mobile advertisers. The mobile advertising world is today not as sophisticated as, say, the online advertising and product-recommendation world. All of this inventive activity is today (summer 2013, the reader may be at a much later point in the development) is closer to creating a new market than fine-tuning a targeting formula.

We should not be surprised at this slow development in light of the historical experience with inventing advertising supported industries. Television broadcasting and receiving were tremendous technical inventions, of course. The economic return to those inventions was dramatically increased by, for example, changes in the nature of certain existing content (e.g. the NFL) the creation of effective means of delivering that content to attract a specific audience (e.g. CBS sports) for adult men, and the creation of advertisers with very large budgets to reach those specific audiences (e.g. Budweiser, a 50% brand, emerging out of a very fragmented US beer industry.) For mobile advertising, those coordinated inventions, which are not (primarily) technical inventions, but commercial ones, still lie largely in the future.

7.3 No (Current) Revenue Stream – corporate apps

Finally, there is a very important class of app which is not paid, not subscription, has no IAP, no freemium, and shows no ads – in short, which has no current revenue stream. Many of these are "corporate" apps, which offer a product or service that is complementary to paid products from a consumer-oriented firm. A banking app, for example, lets a consumer check balances or take a picture of a check for deposit. An airline app similarly lets a traveler display an electronic boarding pass or check the seat map for their flight.

These corporate apps are large and growing as a portion of all app downloads. By the time this chapter is published, the class of app will likely be widely recognized. But after we first discovered it, leading industry figures told us for many months that we were simply mistaken (they have now stopped). These corporate apps represent a change in the mechanism by which the mobile app economy delivers value to the broader economy. The original idea was that apps would offer advertising services to bring in new customers to consumer product and services companies. The corporate apps are typically given by the corporation to existing customers, and offer a wide variety of forms of improved customer service. That, of course, may raise demand, but not (directly) through bringing in new customers.

Monetization of the corporate app is most direct for mobile commerce retailers. Their apps make purchasing opportunities more accessible as impulsive purchasing messages surround consumption driven mobile users. A push notification from a daily deal site could easily induce a consumer riding the subway to make an impulse purchase in a physical environment where monetary transfer may otherwise be limited to giving a beggar spare change. Even without push notifications, apps allow firms to sell tangible goods from locations where traditional e-commerce would never have been possible; eBay has been particularly effective in letting both buyers and sellers access from anywhere. Banks have picked up on this anytime, anywhere theme strongly in their messages to their customers. Without quite saying that banking is an annoying chore, they have suggested one can complete it in times or places where consumers' value of time might otherwise be low.

It is of course entirely possible that the strong turn to corporate apps is temporary; indeed, we argue in this paper that it is driven, to a considerable extent, by the ineffectiveness of the online app stores (corporate apps go to existing customers and escape any problems with app discovery in the app stores.) For now we note, however, that the race between *disruption of everything* and incremental improvements in quite a few things is not being won by disruption.

7.4 Other (currently) zero-revenue apps

Seller monetization strategies may be dynamic. An app developer may therefore choose to have no revenue in the present in anticipation of having revenue in the future. We have talked with many industry participants about this, and four main themes arise. Our interpretation of these themes is that the value of delaying revenue applies particularly to entrepreneurial developers, and within entrepreneurial developers with particular characteristics. Of course, the ability to finance a delay in revenue is also particularly challenging for entrepreneurs who have not been able to secure VC or angel financing.

As a threshold point, in some circumstances an app may be building up a large volume of users with a plan to somehow monetize in the future. Industry participants emphasized emphasize a short list of variants on this theme.

First, an app that is going to be monetized may first want to reach efficient scale. The structure of the online stores means that minimum efficient scale is measured, at least in part, against the rate of downloads. Thus a firm that seeks to have a paid version in the future may have only a free one in the present to avoid sliding up the demand curve. Similarly, the supplier of a advertising supported apps may seek to avoid annoying users in early stages, i.e., show them few or no ads while the app is gaining volume.

Second, some app developer firms may be looking forward to a merger or an IPO and thinking about their (equity) market valuation at that point. Many, many developers and their financial advisers believe that the equity markets have valuation models that depend on user headcount.

Third, at this very early stage of the industry, the ultimate point of writing an app may not be to have a successful an app. For example, the point of the app may be to advertise the development capabilities of the team. Members may hope to be hired into established businesses. The entire app development team may seek to be hired as a group to convert its app into one that creates value post takeover.

Fourth, some entrepreneurial developers simply have no commercial interest at all. Their app solved a problem for them, and they are quite pleased that it solves a similar problem for others. Their bills are paid by their day job. Our industry sources tell us to expect such “hobbyists” more (today) on Android than iOS.

7.5 Some Statistics

In this section, we quickly examine some statistics about app developer monetization strategies. We first look at Figure 5, which comes from an industry source.

Figure 5: Performance of revenue models in mobile apps

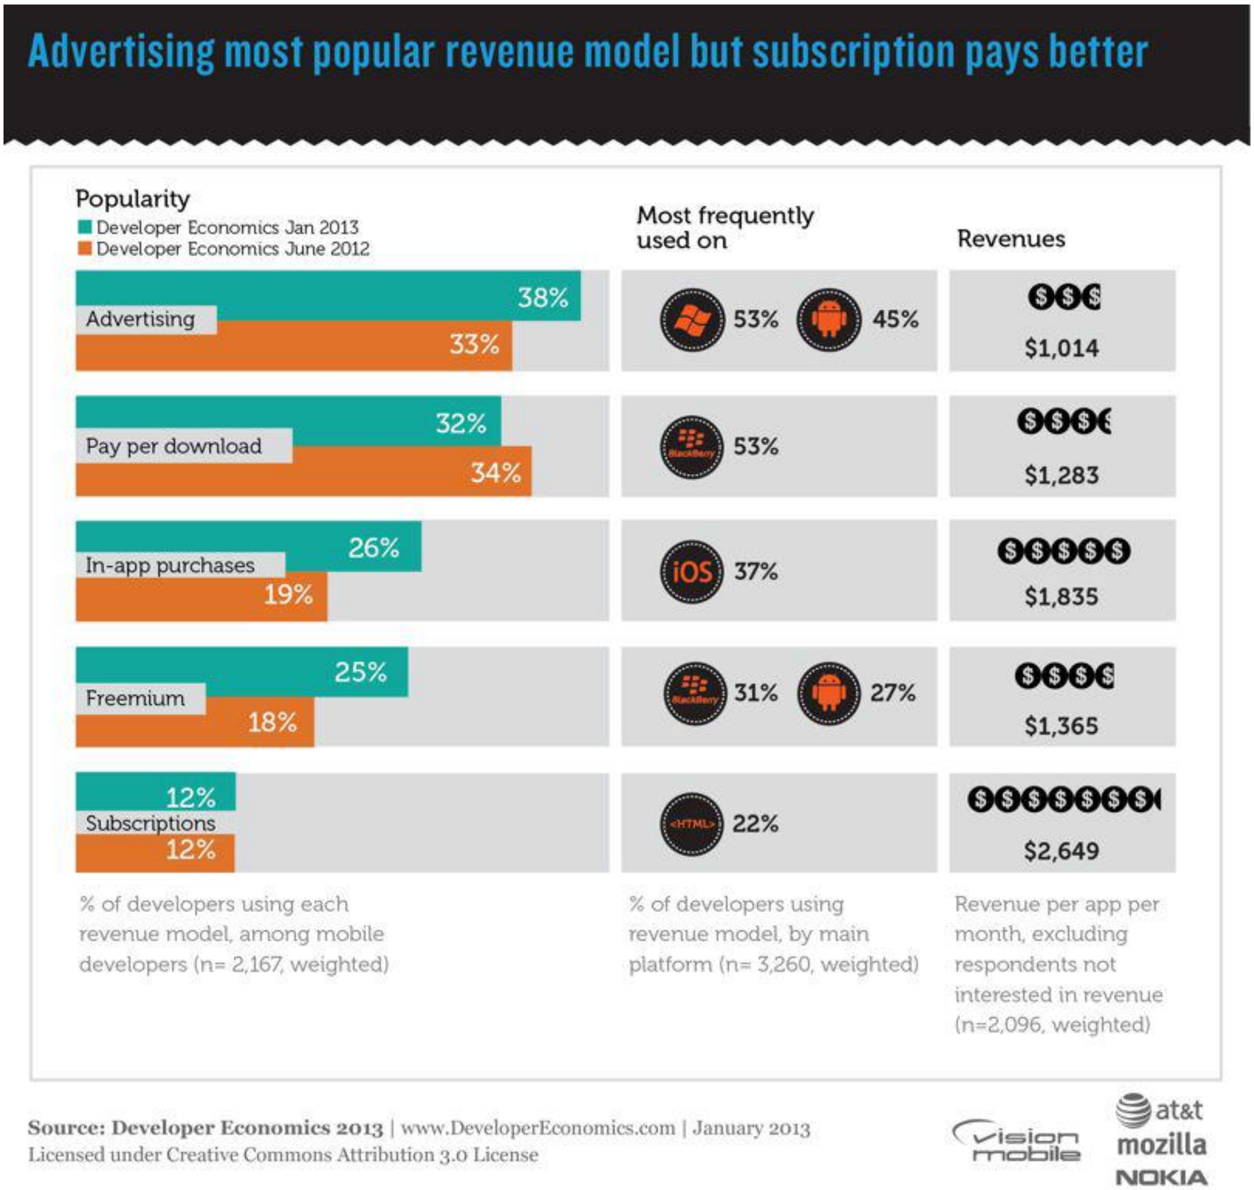

Note that the category of corporate apps is not present, as the Developer Economics table follows current industry practice of focusing on "monetization" as a topic in management-normative analysis.

We can get another estimate of the relative size of the different categories of revenue strategies by examining our app questionnaire data. In Table 4, we show revenue strategies for the most popular free apps on the Android platform, and for a sample of iOS apps which are of comparable popularity. The Table focuses on the joint distribution of any use of advertising and of any use of IAP in free apps. This yields four categories: (i) neither advertising nor in-app purchases; (ii) no advertising but in-app purchases; (iii) some advertising but no in-app purchases; and (iv) some advertising and in-app purchases.

Table 4: Frequency of use of IAP and advertising in free apps

	(i)	(ii)	(iii)	(iv)
Ads used?	No	No	Yes	Yes
IAP used?	No	Yes	No	Yes
	neither	IAP only	Ads only	Both
Android	.394	.0772	.449	.0798
iOS	.221	.108	.394	.277

Source: App Questionnaire. n: Android 2,281, iOS 2,713.

On both platforms, the most widely used monetization strategy for free apps is advertising without in-app purchases, which was used in 2012 for 44.9% of free Android apps and 39.4% of free iOS apps. The most visible difference between the platforms in this table is that in-app purchases are considerably more widely used on iOS compared to Android, especially in conjunction with advertising (27.7% on iOS compared to 7.98% on Android). The higher use of IAP on iOS is consistent with industry surveys (such as the one displayed above as Figure 5) and likely stems from the tendency of iOS users to be richer than Android users.

This table also reveals that more than a third (39.4%) of free Android apps neither use advertising nor in-app purchases to users, and similarly for iOS, about a fifth (22.1%.) This highlights the significance of the class of apps which do not yet pursue a standalone commercialization strategy that is tied to the app economy.[22]

As a robustness check, Table 6 presents the same set of frequencies as Table 5 but for a subsample of the app questionnaire data. Recall that the app questionnaire sampling frame is comprised of the union of two criteria: either an app is identified in the comScore data, or it satisfies the criteria for being a top app on iTunes or Google Play. The subset in Table 6 are the apps which satisfy the intersection of the two criteria: these apps both are top apps and appear in the comScore data. Furthermore, the subset is limited to those apps in this intersection which actually had been processed by our research assistants as of January 2013. The result was 431 free apps over both platforms. Note that despite this more limited sampling frame, we still get the same results as we do in Table 5.

Table 6: Frequency of use of IAP and Advertising in free apps

N=431	(i)	(ii)	(iii)	(iv)
Ads used?	No	No	Yes	Yes
IAP used?	No	Yes	No	Yes
	neither	IAP only	Ads only	Both
Android & iOS	.258	.135	.357	.251

Source: App Questionnaire as of Jan 2013 restricted to intersection of comScore & top apps

[22] It is possible that the lower incidence of these apps on the iOS side is an artifact of our sampling frame for apps, which, because of the ITunes Store's higher-frequency collaborative filter, may have picked up more apps that were "buying" distribution. These apps would tend systematically, *not* to be corporate apps. We are investigating this using an alternative sample based on ComScore.

In Table 7 (in progress), we further subdivide the apps within each revenue model into those associated with a non-mobile consumer products and services (pre-existing, "corporate") firm and those associated with a purely mobile firm. This is a limited subsample, since the question of firm type was not asked in recently processed surveys. We interpret the former as monetization by improved customer service, the latter as the early stages of a dynamic monetization strategy. The "corporate" apps (are expected to) dominate the category of no-current-revenue-stream apps.

7.6 Advertising demand by Mobile Apps

We have already noted that the collaborative filter of the online apps stores is an important barrier facing apps that do not have an existing connection to consumers. This gives app developers an incentive to build volume quickly. This can lead to the irony, (now) frequently noted by industry observers, of an app which plans to be supported by advertising in the future but which, today, has zero revenue but spends on mobile advertising to get more downloads. More generally, any app without an external body of customers has some incentive to keep up its rate of downloads at all times, even at those times when that might be difficult, such as right after introduction, just before a significant upgrade, and so on. App developers are demanders of advertising space in apps. One way for app developers who have a budget to advertise is to buy space through an ad broker. There are also ad-exchange clubs, in which app developers agree to show one anothers' ads for apps. Finally, multiproduct app developers can run an ad for their own products, either another app, an in app purchase of a virtual good or some other kind of product.

Above and beyond advertising to build volume, app developers can pay to have users download their app. These are called "incentivized" downloads in the industry, and a thriving and changing business has grown up to supply them. Firms such as Tapjoy have flourished in this space. Suppose app firm A has a successful app with a virtual currency used for IAP. App firm B wants downloads. Then B pays Tapjoy to pay A to offer virtual currency to A's customers in exchange for downloading B's app. Industry sources tell us the strategy of buying users tends to be more prevalent in apps that monetize through the sale of virtual goods, subscriptions or apps that are an arm of an existing consumer-oriented commerce firm. It appears that the strategy of buying users is closely related to those apps which have an anticipated high average revenue per user.

Like advertising focused only on mass downloads, "incentivized" downloads may not lead to users who are good customers (i.e., who use the app and monetize well.) This is a recurring complaint among industry sources we interviewed. This conundrum has created further innovations, including firms that incentivize users to watch movie-style trailers of apps, letting the user decide if she should download the app, and firms that incentivize but give the user a choice of apps to download. We anticipate that ad-tech firms will continue to innovate new ways to "buy" users in the coming years, potentially discovering new solutions to incentivize users with a high likelihood using apps and spending money inside apps.

We can examine some of this behavior of "buying users" by looking at the identity of advertisers. To do this, we return to our app questionnaire data. We asked students to download apps and answer questions about both the location and the content of the ads.[23] We count all the different ads that show up in one use of the app – defined as, at a minimum, passing all the places where ads might be shown. Thus the probabilities sum to (considerably) more than one. The

[23] To the extent that there is targeting, the resulting sample of ads is composed of either mass-market ads trying to hit more or less all users, ads targeted to people who use their phones like students, or ads targeted to people who download many apps (since our students download a significant number of apps to fill in our survey.)

dominant location is banner ads, so we report the results only for those. This exercise treats paid and unpaid (same firm, or bartered through an ad exchange) the same.

Table 8: Banner ad content

	Android	iOS
Apps with banner ads	N=1106	N=1583
any app ad.	.802	.847
any non-app ad.	.370	.371

Source: App Questionnaire

Table 9: Percent of apps with banner ads in each category (weighted average over platform)

	App ad	Non-App ad
Same Firm	0.348	0.091
Different Firm	0.623	0.295

Source: App Questionnaire

A surprising finding from that Table 8 is that if an app has a banner ad, there is an 80% chance or greater that it will serve an ad for an app! In Table 9 we see that there is quite a bit of self-advertising in this industry. Over a third of the apps with banner ads advertise their own apps. The proportion of apps with banner ads that advertise other firms' products outside of the mobile app industry is quite low at under 30%. The demand for ad space in apps is demand for space to advertise other mobile apps.[24]

In Table 10, we report four categories of ads, based on the kind of product advertised and the identity of the advertiser. For this purpose we have only two kinds of products, mobile apps and all others, and only two identities for the advertiser, the owner of the app in which the ad is running and all other. We recorded data for more categories than this, but once you see the table it will be obvious why this is what we report.

Table 10: Advertised Products and Vertical Integration in advertising

	Android Banner Ad Apps n=1106		iOs Banner Ad Apps n=1153	
	App ad	Non-App ad	App ad	Non-App ad
Same Firm	0.102	0.0171	0.486	0.133
Different Firm	0.727	0.357	0.565	0.26

Source: App Questionnaire

This is a complex enough table that it is useful to walk through the numbers. Table 10 reports the products and advertisers seen in Android and iOS platform apps that displayed at least one banner ad. Within the Android banner apps, about ten percent (.102) display a banner ad for

[24] When we first talked with app-oriented venture capitalists about this fact in late summer 2012, many thought it was simply false. Today (summer 2013) savvy VC accept the necessity of buying downloads.

another app from the same firm, while just about 2% (.017) display a banner ad for another product from the same firm that is not an app. An example of the latter would be a media company advertising a television show in its app, a common ESPN behavior for instance. Continuing with Table 10, over 70 percent of the Android apps (.727) that have a banner ad have a banner ad for an app from a different firm. Finally, just a bit over a third of these apps displaying banner ads on Android (.357) display an ad for a non-app product or service from a different firm.

In contrast, we find that for iOS, about fifty percent (.486) display a banner ad for another app from the same firm, while over 10% (.133) display a banner ad for another product from the same firm that is not an app. Only 57% of the iOS apps (.565) that have a banner ad have a banner ad for an app from a different firm. Finally, just a bit over a quarter of these apps displaying banner ads on iOS (.26) display an ad for a non-app product or service from a different firm.

Self-advertising is much more important on iOS because 1) a developer can more easily influence an apps placement in the rankings on iTunes relative to Google Play since the rankings only use the last 24 hours rather than the last week of downloads, and 2) since iTunes is the only distribution channel for iOS apps, influencing iTunes rankings is relatively more important to an apps success than Google Play rankings (which marginally competes with alternative Android app markets). Apps with banner ads advertising a different firm's apps are less likely on iOS since those apps might represent direct competition in the rankings. Note that since the percentages do not have to add up to one, there is no mechanical reason why self-promoting app ads would be higher on iOS and competitor-promoting app ads would be lower on iOS.

These facts come from an early stage in the development of the industry, and they also come from our students' phones, i.e., not necessarily from the most valuable advertising audience. Given that potential oddity of the sampling frame, we think that there is a strong conclusion and a weak conclusion. The strong conclusion is robust to our sampling frame: app developers today have a powerful motivation to buy downloads in order to become visible on the online app lists.

This conclusion is reinforced by the different behavior of corporate app developers (who have that motivation much less) and entrepreneurial app developers (who have it much more) since these two groups of developers have very different exposure to the costs associated with the collaborative filter of the online app store.

Our weaker conclusion is that exports from the entire sector of mobile app advertising to the rest of the economy are growing slowly. Here our sampling frame may matter, to some degree. We want to point out that robust revenue numbers for app advertising don't rebut this finding. Those numbers come from summing the revenues across app developers,[25] without netting out the within-sector sales.

The rate and direction of application innovation is being affected by the need for developers to devote resources to solving the matching problem and getting noticed out of clutter, rather than devoting resources to monetization efforts based on creating value for customers and rather than trying to gain money from marketing products outside of apps.

8. DEVELOPER BEHAVIOR: PLATFORM CHOICE AND MULTIHOMING

In this section, we document the multi-homing behavior of developers on iTunes and Google play. We pick a good moment for a snapshot (the platforms are balanced now and expectations don't seem to tilt) and (2) a study of how Android caught up is within the data. We show that there is relatively little multihoming, and multi-homers are characterized by existing

[25] "Revenues" in the industry sources are themselves not the strongest numbers in the economy.

capabilities and assets which allow them to solve the matching problem to their customers outside of the mobile apps world. We view this as evidence that the matching problem slows the rate of multihoming for most apps, and therefore slows the rate of application innovation and the resolution of the platform competition for mobile application supply.

At present, an app developer can reach approximately half of all mobile users by writing only for iOS and distributing through the iTunes stores, and approximately half of all mobile users by writing only for Android and distributing through the Play Store, Amazon, and so on. A much smaller user body can be reached by writing for some other platforms, such as Windows mobile, Blackberry, and so on.

There is a lively debate in the industry about how the two main platforms differ and about ideal strategies for a developer in choosing among them. The debate emphasizes the asymmetries between the two platforms, summarized in Table 11.

Table 11: Asymmetries between platforms

iOS & iTunes Store	Android & Google
Early: more devices in use	Now: caught up in total devices in use (mostly phones)
More Tablets	Tablets rapidly growing
More commercial infrastructure including payment processing tools	Absent, but catching up
Richer Users,	Less rich users …. may not buy IAP, for example
More Restrictions on Developers	
Development Environment on Macintosh	Develop anywhere
Apple dictates tools to be used (e.g., flash)	Use Java, popular with developers
Limited range of devices	Fragmentation
"Managed" change from year to year … porting an app to the newest iPhone/iPad devices from older ones usually simple	Changes in environment from year to year, e.g. substantial UI changes Different hardware manufacturers use different OS versions
Distribution restricted to iTunes	Open; multiple distribution channels

Matching apps across the two platforms is significantly more difficult than you might at first expect. The same app can have different names in the iTunes store and Google Play, and different apps can have the same name. The supplier has only one name on Google Play, and that can be the same as any or none of the three supplier names on the iTunes store.[26] Our best results have come from looking at developer websites one by one and capturing all the apps on both

[26] Computer scientists have built sophisticated name-matching software for this and other purposes (and enthusiastically recommend it to us regularly.) The best software solutions for matching in this context correctly identify just under a half of true matches – that is, have a more false negatives than true positives – and also identify a significant number of false positives, about one tenth of the number of true positives.

platforms (where there are unique identifier app ids) claimed by the developer.[27] We employ our app questionnaire and the comScore to examine the current state of multihoming.

We begin with the January 2013 comScore sample.[28] In that month, there were 1,231 apps in the comScore data set. In Figure 6, we show the distribution of apps among iOS only, Android only, and Both categories. The low rate of multihoming is something of a mystery, given the large market available to apps on the platform they do not serve, and that we are examining the most popular apps. A disadvantage of the comScore for this analysis is that the coverage of apps is narrow since comScore measures presence on devices, not availability (so that a developer who multihomes but has limited market success on one platform may be missed).

Figure 6: Multihoming, iOS exclusive, and Android exclusive app percentages

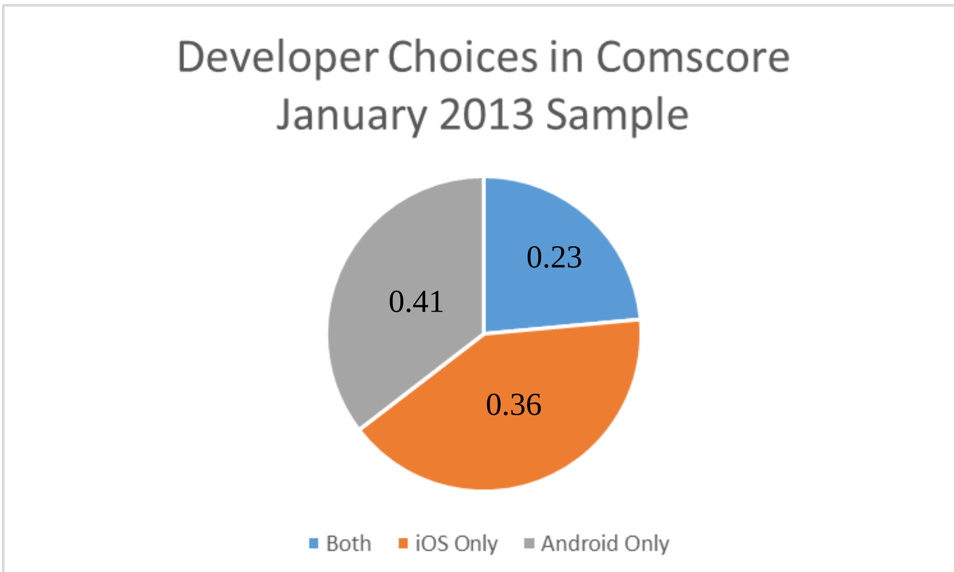

Source: comScore

We noted previously that the size distribution of app demand was highly skewed. The larger is the demand for an app, the more likely it is to be profitable for the developer to bear the incremental fixed costs of multihoming, i.e., the technical costs of porting to the second platform and the marketing costs of establishing a connection to consumers on the second platform.[29] That would imply a higher rate of multihoming for more popular apps, a hypothesis that can be examined by looking at different *weighted* propensities to multihome.

In Figures 7a and 7b we display the rate of multihoming under five different weighting schemes.[30] Figure 7a shows the rate of multihoming and single homing for apps which are observed on iOS: unweighted (the last bar), 60% of apps found on iOS are <u>not</u> found on Android.

---

[27] When we are working with only the ComScore data, we can use ComScore's matching. But ComScore's names for apps differ from both the ITunes store and Google Play. Accordingly, we have also hand-matched ComScore "apps" to unique identifier app ids. We very much appreciate the help ComScore has given us in resolving the last few difficult cases.

[28] The ComScore mobile product is new, and had an initial period of rapid improvements in coverage; January 2013 is after the ComScore mobile sample and definitions settled down.

[29] Cite Bresnahan and Reiss (19xx).

[30] DOCUMENTATION!!!!!!!!!!!!!!!!!! The schemes are daily "visitors" (i.e., users), minutes per "visitor", i.e. the average amount of time the app is open when it is being used, total minutes (over all users), unique "visitors" over the month, and unweighted. Definitions of all these different weights are found at XXX.

Figure 7b, symmetrically, shows the rate of multihoming and single homing for apps which are observed on Android. The first result can be seen by comparing the last bar (unweighted) in either figure to the bar next to it (weighted by unique visitors). In each figure, we see that the weighting reduces the rate of single homing and increases the rate of multihoming, exactly as one would expect if apps with a higher market size are more likely to multihome. Very similar results arise comparing the last bar (unweighted) to either bars 1, 3, or 4; this, too, is unsurprising, since those three weighting schemes are based on very similar measures of an app's market size. The effects are not small: weighting by any of the measures of market size (e.g. number of average daily visitors, total minutes per month, and number of unique visitors) then the percentage of multihoming exceeds 60% for both platforms. Nonetheless, even with this weighting, the incidence of multihoming is surprisingly low. We conclude that the predicted market size threshold for porting to a second platform is high. Of course, that threshold varies with many other variables, so we don't have solid estimates of it from these simple investigations. Nonetheless, our sample of apps, small as it is, appears to read far enough down into the "long tail" of apps that for many of them it is not economic to port.

The second bar ("minVisitor") in each of these figures uses a very different weighting scheme as average minutes per user. This measure is related to the intensity of preference at the individual level – typically called "user engagement" in the industry – more than to the size of the market (number of users) for a particular app. This weighting scheme is quite similar to the unweighted bar – a bit higher on iOS and a bit lower on Android – and quite unlike the market size weights. It does not appear that the apps which are more likely to engage their users for a longer period of time are more likely to multihome. There are two obvious explanations for this, and we are not yet able to distinguish them. First, high-engagement apps may be more difficult to "port" from one platform to another, either because it is harder technically to write them for the other platform (e.g., they may be games with much use of the user interface) or because it is more expensive to find an audience for them. Second, having the same user for twice as many minutes may not be as close to getting twice as much profit as is having twice as many users. This would arise for paid apps, or because high engagement apps cannot show advertisements (profitably) twice as often if the app is open twice as long.

Figure 7: Multihoming of apps for different weighting schemes, January 2013

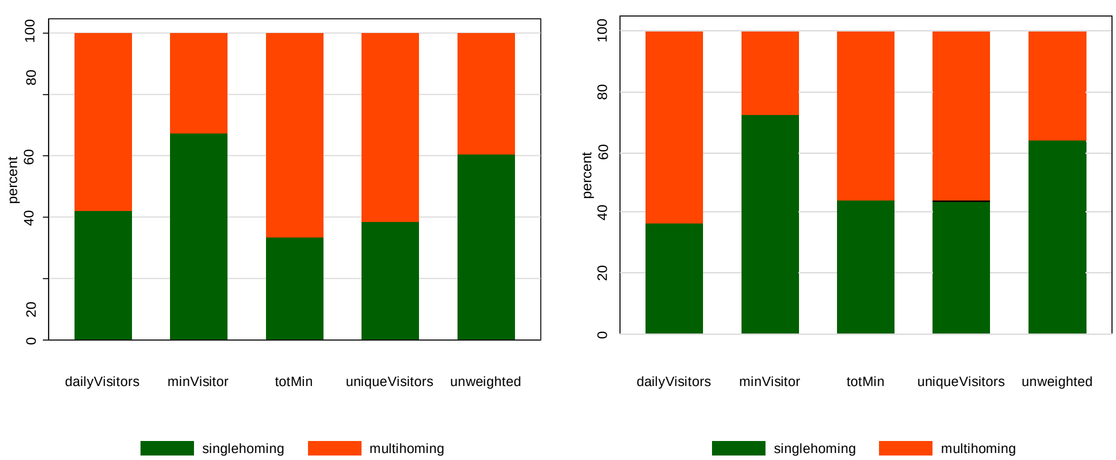

(a) iOS apps (b) Android apps

The slight preponderance of Android apps over iOS apps in these data might reflect any of several things. It is important to recall that the criteria for being included in comScore are a very low cutoff for reach – about one in a thousand users. We saw above that the size distribution of apps is highly concentrated. A one in a thousand users cutoff takes us well out into the "long tail" of very small app design. Thus the mere count of apps might, on either platform, include a fairly large number of marginal apps. Since the costs of entry on the Android side are a bit smaller, and since many developers of apps without strong commercial capabilities also have an open system preference, there could be a larger number of commercially larger apps on the Android side. Second, the relatively small size of the comScore mobile panel means that apps whose true reach in the broad population is near the cutoff (.001 reach) will only be included in the data probabilistically. Without careful investigation of that ordinary sampling uncertainty, it is not obvious that the difference between the two platforms in app counts differs statistically from zero. Both of these points leave us with a weak conclusion: it is not tipping to iOS, but the foundation for an Android advantage at this stage is weak.

We can learn a great deal more about the economic motivations for single-homing by also examining our second app questionnaire dataset.

In the app questionnaire dataset, our definition of multihoming is that the developer *offers* the app for both platforms. Not surprisingly, this yields a higher rate of multihoming than the comScore data, in which we observe multihoming only if the app is offered on both platforms and *demanded* by .0001 or more of users on both platforms. Looking only at the same apps, we observe 23% multihoming in the comScore data (Figure 6) and 28% multihoming (Figure 8) in our app questionnaires. The gap between app *offerings* on the second platform and app *demand* is, we think, a clue that the barriers to porting to a second platform are not merely the technical costs of rewriting the app, but also the marketing costs of finding an audience.[31]

[31] Of course, a more careful analysis of the exact nature of the marketing costs of finding an audience on the second platform and of the impact of those costs on supply would be necessary to make this inference complete. An important part of such an analysis would be to distinguish between app quality and marketing costs as alternative theories. Sorensen (20xx) has an interesting analysis of precisely such a problem in a related context, "best-seller" lists. The problem of finding an audience for an app on a platform is, to a considerable degree, one of breaking into a "best seller" list.

Figure 8: Multihoing, iOS exclusive, and Android exclusive app percentages

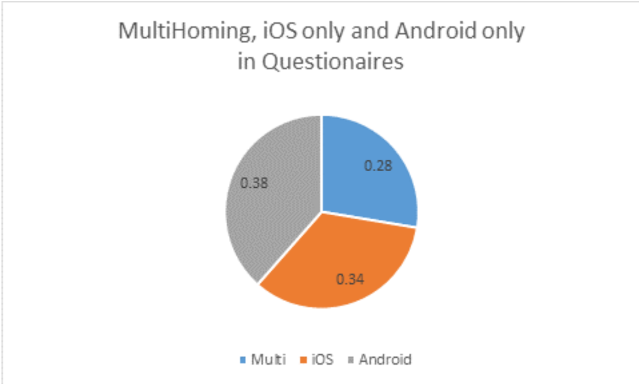

Source: App Questionnaires

A second clue to the same point can be seen in a limited firm-level analysis shown in Figures 9a & 9b. In these figures, we consider the top apps drawn for our App Questionnaire from iOS separately from the top apps drawn for our App Questionnaire from Android. The developer's website was accessed (when available) for each app, and the questionnaire determined whether the website claimed that the focal app was offered on the other platform. The questionnaire also asked whether the website claimed that any other apps, even if not the focal app, were offered on the other platform. Since developer websites were not always available, the data is limited to a subset of the App Questionnaire apps. Figure 9a shows that developers of 23% of successful apps on Android will claim to have ported that app to iOS. Successful iOS apps have 29% probability of being ported to iOS. However, Figure 9b shows that sometimes developers multihome in that they develop apps for both platforms, but they offer different apps for different platforms. Figure 9b shows that there is a 29% probability that the developer of a successful app on Android will claim to offer at least *some* app on iOS. A developer with a successful app on iOS has a whopping 49% probability of claiming to offer some app on Android. The difference between the values in Figures 9a and 9b suggest that there are a number of apps whose developers already multihome and know how to develop for the other platform, yet choose not to port all apps to the other platform. This is evidence that the primary blockage to multihoming in these cases is not the technical cost of porting.

Figure 9: Developer website claims that they develop apps for the other platform

9a) this app

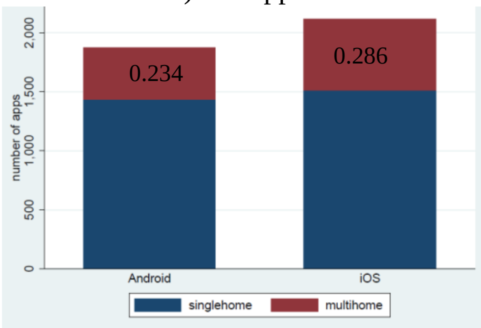

9b) some app

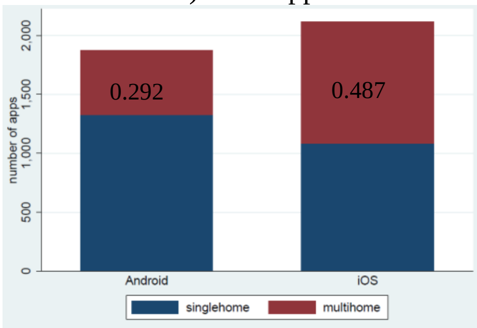

Source: App Questionnaire, n: Android=1,874, iOS=2,117

In the app questionnaire dataset, we have also constructed variables for the type of firm associated with the app developer (purely mobile app firm or a pre-existing, “corporate” firm offering non-mobile products). We report the rate of single homing and multihoming for each observable type of firm in Figure 10. Since this question was not asked in recent versions of our survey, we are limited to a subset of the apps in the App Questionnaire dataset.

Figure 10: Multihoming by firm types

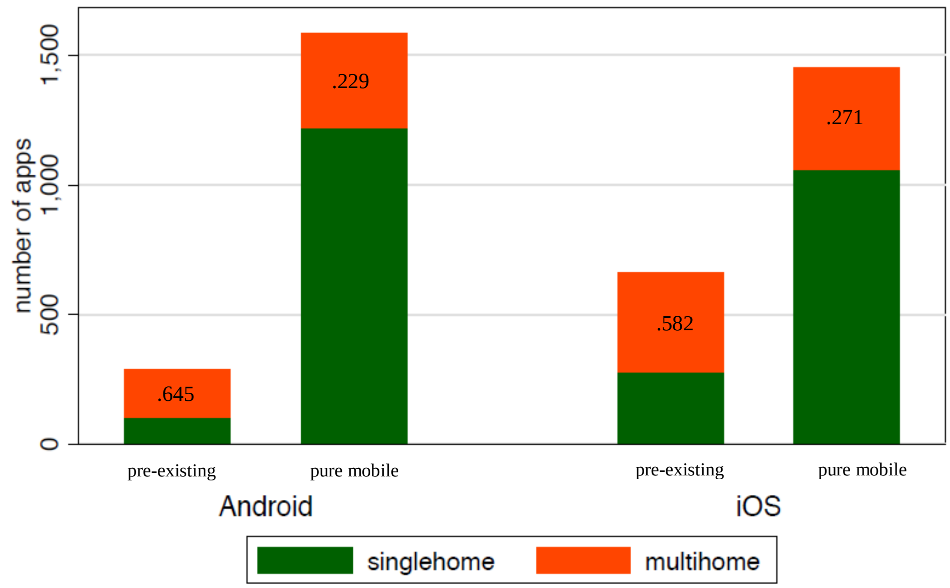

Source: App Questionnaire, n: Android=1,874, iOS=2,117

Figure 10 also provides considerable evidence for the hypothesis that the costs of porting and multihoming are comprised, in no small part, of the marketing costs of finding an audience on a second platform. Top Android apps whose developers have a connection to their customers from outside the mobile world, such as existing consumer product and services companies, companies with an online presence, etc. are 65% (.645) more likely to multihome. For iOS, the probability of a top app multihoming if the developer is pre-existing is 58% (0.58). These probabilities are 2-3 times higher than the probabilities for apps from purely mobile developers. Such firms have no need to find a new audience on the new platform, and thus have far lower incremental marketing costs of multihoming.

There are a number of limitations to this discussion. The difficulties in matching apps across datasets means that we do not have a fully integrated dataset yet. Furthermore, we have not filtered out apps for whom the porting decision may not be interesting or may already be pre-determined (wireless carriers, platform providers, etc.). A great deal of interesting analysis may arise from augmenting this data with information on categories (c.f. Davis et al., 2013). Finally, we as yet do not have a serious model for understanding first and second platform choice by developers.

8.1 Competition among Platforms

The topic of the implications of all this developer behavior, including single-homing, multihoming, and platform choice for the single-homers, for the competition between platforms is a difficult one at this juncture. No tip to a single platform seems immanent, even though multihoming is rare. The Android/Google Play and iOS/iTunes Store platforms are approximately equally attractive to developers today, which suggests temporary stasis near the unstable platform market equilibrium of approximately even market shares for the competing platforms.

Modern platform competition theory emphasizes the critical role played by expectations in resolving which of the many platform market equilibria could arise. [32] There does not appear to be a strong expectation -- in any direction -- among consumers.
To the extent developers have expectations, those would appear to favor the original incumbent iOS/iTunes Store platform. (See Table above.) Money-oriented developers are directed toward the richer customers of the iOS/iTunes Store platform, even today, and of course for a while that platform had a strong lead in installed base in smartphones (and then, for another while, a strong lead in installed base in tablets.)

However, many developers appear to be treating the two platforms as broadly equal. Large corporate developers tend to write for both platforms. Also writing for Blackberry or Windows Mobile is much rarer for them; our interpretation is that they view the two largest platforms as much more valuable than the new smaller ones. Smaller developers are much more likely to be single homed. As we saw above in Figure xxx [pie chart] the count of smaller developers in the comScore data with Android-only exceeds the count with iOS-only. For now, our interpretation is that the market is <u>not</u> tipping to early leader iOS. Whether a more careful investigation of developer behavior would reveal a tilt of those apps more influential on users toward either platform remains to be seen.

Other entertainment uses of smartphones and tablets have a more multihomed structure today than do apps. First off, the Amazon Kindle platform and the related Amazon Kindle online store represent an important media-distribution and media-reading system, certainly a third alternative to iPads and Android tablets. Second, media (books, newspapers, music, etc.) are significantly more likely to multihome than are app developers. Some of that multihoming is partial, as when iPad users can run an Amazon Kindle app in order to read books they have obtained from the Kindle store. Still, there are powerful forces toward concentration in the media-reader function, just as there are in the app-platform function. These forces were demonstrated by the recent exit of the Barnes and Noble Nook.

Whether there will continue to be significant differentiation between special-purpose media reader devices (Kindle Fire) and general purpose mobile devices which also support media reading (iPad, Galaxy) remains to be seen. This is a second question parallel to, and analytically very similar to, the question of a platform “tip” within the applications platform function.

In short, we are in the “before” period in which a large number of alternative solutions to the same problem are racing in the marketplace. Whether the ultimate important platform(s) will be open (Android/Play Store) or proprietary (iOS/iTunes Store), whether they will be general purpose devices like those or (also?) special purpose media readers, all remain to be seen. So to, does the boundary between all these mobile devices – today primarily supporting consumption – and PCs – today primarily supporting work – in market and in functions served remain to be seen. Perhaps the most important development so far in the mobile devices running iOS or Android is that they have drawn enough users and app developers to create significant momentum, so that there is a real question not only about the competition for ultimate leadership in what is now the “mobile space” but also about competition between these platforms and existing e-reader and PC platforms.

[32] See, e.g., Farrell and Klemperer (20xx). This handbook chapter particularly focuses on the distinction between expectations that track efficiency (developers and users expect the better platform to be adopted) and inertial expectations.

8.2 Divided Technical Leadership

It is a fact of information technology industry life (if not a feature of "two sided market" models) that many application platforms partially or sometimes fully overlap with one another. Today, for example, there is a positive feedback loop around not only the iOS/iTunes Store platform, but also around the Facebook applications development platform and around Google Maps – while both a Facebook app and a Google Maps app run on iOS devices. This structure, which causes divided technical and market leadership of the positive feedback loop in any one of the partially overlapping development platforms, has implications for the analysis of platform competition.[33]

One implication of divided technical and market leadership is that it can make change the nature of competition between substitute platforms. Entry and competition against the established platform of IBM mainframes was, for example, greatly facilitated by the existence of the partially overlapping platform of database management systems, and firms such as Oracle were important complementors of entrants competing against the established IBM platform. A second implication of divided technical and market leadership is that it weakens, sometimes dramatically, the control of a platform sponsor over technical progress within a platform, permitting innovation competition among the different firms sponsoring the different layers. The influence of firms other than IBM on innovation in the "IBM PC" is a famous example.

The Android/Google Play mobile platform was set up by sponsor Google as if to enable divided technical and market leadership. Android can be altered by other firms, and has been, for example to create e-readers which are now outside Google's control, to create special versions of Android for particular manufacturers, for particular screen sizes, and so on. Similarly, app distribution could flow through alternatives to Google Play, such as Amazon. Accordingly, Google cannot prevent the widespread distribution of apps, and thus permits app developers wide leeway. There are, as a result, a number of Android apps that offering "infrastructural" functionality, functionality that many operating system providers would seek to offer only in the OS itself.

Apple set up the iOS/iTunes Store platform in a much more centralized and controlled way, and thus can block the emergence of apps it feels are bad for consumers or not in its own interest. As a result, divided technical and market leadership on the iOS/iTunes Store platform cannot emerge easily. However, a number of partially overlapping platforms, notably from Facebook and Google Maps, have created some divided technical and market leadership for Apple. Apple's unimpressive effort to replace Google Maps reflects the value to a proprietary platform sponsor of preventing divided technical and market leadership from emerging.
One significant influence on the direction of the mobile platform competition among Google, Apple, and potentially Microsoft, Blackberry (nee RIM) or others arises from highly influential suppliers of partially overlapping platforms. Should, for example, Facebook's influence with mobile users grow rapidly, that firm could have considerable influence on the market equilibrium for mobile operating systems.

9. ALTERNATIVE EQUILIBRIUM SCENARIOS

[33] Bresnahan and Greenstein (19xx) argue that this structure has been historically important in competitive parts of the computer industry, such as the early PC industry. The structure has been important both in enabling new platforms to compete against and established platform and in enabling competition within a platform even with positive feedback around a particular standard.

While the current platform industry structure seems highly unlikely to persist for a long period of time, it is less than obvious what ultimate market structure will arise. Obviously a tip to either iOS/iTunes store or to Android/Google Play could occur. Less likely, one of the newer platforms (Windows mobile, Blackberry) could get over the network effects hump and the market could tip to it. Given Microsoft's position in PCs, it seems certain that that firm will continue to invest in Windows Mobile for a long time (a la Bing) even if it cannot move beyond a second place platform. Another potentially important long run equilibrium scenario is platform product differentiation. Many industry participants imaging a future in which Android is stronger in selling to poorer customers than iOS, possibly poorer customers in the rich countries or possibly becoming the dominant platform in poorer countries. Like a platform tip, a platform product differentiation equilibrium involves a substantial change in app developer behavior. At this stage, we do not see any strong precursors of any of these long run scenarios. The average market participant who comments on this issue suggests a tip to Apple, a scenario that seemed to us more likely a few years ago than it does today.

Finally, the OS and online store platforms we see today could be commodified by the movement of control of standards to a different layer in the value-creation "stack." Control of mobile development standards by Facebook is one possible scenario here, with the authors of the mobile device OS becoming less influential on developers (a la the shift from "the IBM PC and clones" to "the Windows PC".)

10. CONCLUSION

We have been writing, for much of the last few sections, about the institutional and conceptual bottlenecks to successful exploitation of a tremendous technical and market opportunity. This is interesting for two reasons. First, it illuminates an important general point about the creation of economic value out of technical progress; the "last step" in the chain of invention that leads from new basic science to new economic growth can be quite difficult, and can, even though it is "merely" the discovery of the most valuable uses for a new technology and the creation of markets to serve those uses, call for profound innovation itself.

Second, the sheer size of the mobile opportunity illustrates an important change of 21st century innovation from earlier rounds of innovation, at least in ICT. There are three quarters of a million apps on each of the major mobile platforms: A lot of experimentation fits into three quarters of a million new products. There has been much success at the firm level, but there is still an opportunity for some of those experiments to create a big hit, finding important new value for mobile applications. At the stage when the PC industry found its first big hits – white-collar work apps like the spreadsheet and the word processor – there were far fewer app innovations (in the hundreds) and a vastly smaller (<1million) device market. We are seeing the benefits of cumulative technical progress joined to the benefits of successful exploitation of social scale economies through the platform model of organization.

Perhaps more interestingly, the very large potential markets created by the rapid growth of mobile device usage, and the low costs of entry created by the platform technologies are an invitation to potential app inventors with a wide variety of knowledge bases and a wide variety of incentives. The resulting entry, even at this early stage, has been heterogeneous as well as numerous. The creation of a global communication infrastructure has meant that the sources of supply of apps have also been global. We can once again draw the analogy to the early PC industry, in which international app supply meant "localization," not drawing on a global talent pool; there was no Rovio, no Distimo.[34]

Similarly, high device penetration has made it economically feasible in the mobile era to have app experiments are very heterogeneous with regard to the developer's business model. What Hayek called the economic function of the entrepreneur (19xx), finding overlaps between supply opportunity and demand need, can be taken up by any organizational form – and has been. Right now, the online app-discovery mechanisms tend to favor existing businesses which give their existing customers an app. That is, entrepreneurial function is being taken up by firms which do not take the entrepreneurial form. But there is no strong reason to believe that this is a permanent situation; the economic returns to the creation of new app-discovery mechanisms are simply too large to believe there will not be new invention in that area.

While many industry participants talk about universal "disruption" of existing economic institutions and markets resulting from mobile systems development, that is largely in the future. A number of very interesting experiments are being tried, and a number of potentially important applications areas are entering a very early stage of a long diffusion process. Industry participants routinely speak as if they know the results of the experiments and as if everyone is already doing what only the earliest adopters are trying out. This is a familiar situation in the applications of information technology, which economists labelled "the problem of the tenses" three decades ago.[35]

Why has the app exploration taken so long? Whoops, there we go, falling into "the problem of the tenses!" Better: Why will app exploration take a long time to create all the markets and other institutions this industry will need for long term value creation? Consider a very simple consumer problem, one where mobile devices and apps have already created a lot of value, of finding a coffee shop in an unfamiliar neighborhood. "Maps" is a good application for this, and it is possible that there will continue to be a very large market in which maps are free to the consumer and (some) retailers pay for ads to be displayed in maps or some other product supplied jointly with maps. Or, since coffee shops are consumed socially, either a general purpose social network app or a special purpose meet-and-greet social app, such as those used frequently by young people to find a bar, could guide consumers to a coffee shop their friends like, one their friends are in, one where potential new friends are sitting, etc. Those different solutions are all technically feasible, and they also all entail different visions of consumer behavior and retailer behavior as well as of app development. Without experimentation, do we really know what, other than caffeine, consumers will want from the system that helps them find a coffee shop?[36] Even more difficult to foresee without experimentation is the market

[34] There were, of course, exceptions to the marked US-ness of the early PC industry, but they had no important influence on the industry's technical or (other than enabling global sales through "localization") market development. See Bresnahan and Malerba (xxx) for a discussion of national vs. international supply in the creation of a number of earlier computer markets.

[35] As we write, Google folks face the difficult problem of continuing to talk about what everyone "is" doing with their Android phones and tablets while moving on to talking about what everyone "is" doing with Google Glass. This is "the problem of the tenses" in the present.

[36] Some grumpy observers have already said that the high weight that app development so far has put on social features (the coffee shop where your friends are or where new friends might be found) arises from 20-something app developers thinking about the concerns of 20-something customers, in this case the incomprehensible tribal and mating behaviors of 20-somethings. This is uncharitable; market experiments are heterogeneous because different experimenters have different knowledge, goals, powers of conjecture. Heterogeneity is very good in markets as large as the mobile industry; ultimate economic value creation does not turn on how many experiments are wrong, but on whether any are right.

equilibrium balance between consumer interests (do I find the coffee shop I like, do shops compete for my business) and retailer interests (do the ads bring me customers I would not otherwise have had.) The (non-mobile) online world continues to explore these equilibrium questions almost two decades on; that pace is determined by the pace of market exploration, not by the pace of technical change.

To continue to use the very simple coffee shop example just for a moment, there is another possibility, which is that rather than a general app like maps or facebook or a special third-party app like Coffee Shop Finder that helps the consumer actively search, a retailer app comes to play a very important role in this area, tying the consumer more tightly to a particular chain of coffee shops. Starbucks, for example, implements a significantly more successful volume discount (loyalty program) through its mobile app than it ever could through prepaid discount cards, and reports that its customers are much more "sticky" when they use the mobile app to pay. They also report that over 10% of their (US) sales are now (early summer 2013) paid by mobile app, which might lead you to think that (some) retailers are not entirely powerless in the struggle with Google, Apple, Facebook, and the vast herd of entrepreneurs over who will get the rents from the mobile opportunity.

There is no particular reason to think that this struggle over rents will play out the same way in all markets; an enormous market-creation cluster of parallel experiments in a large number of consumer markets awaits. There is no particular reason to believe that the momentary advantage given to large, pre-existing firms (like Starbucks) by the app discovery process today will persist. That too, as we have said, could be changed by new innovation. What there is every reason to believe is that the incentives for new commercial innovation over the next decade created by the opportunity – incentives for a wide variety of new and existing firms – are enormous.

We are not arguing that there is no widespread prospect for disruption, rather the reverse. The mobile developments have already impact in considerable impact on a few areas. Music and other media, already going through a dramatic change because of the Internet, see that accelerated by mobile. (There is likely more to come, as firms like Spotify, Last.fm, Pandora are in competition with the music portion of the online stores.) Mobile telephone carriers have found their business radically changed, and are (mirabile dictu) embracing open systems at long last. But many of the other obvious loci for disruption from mobile are still in the future. For example, radical change in advertising markets is still to come.

We have written much about the bottlenecks holding back this progress. But it is clear that the market process is working rapidly to resolve these problems. Already, we see a tremendous market response to the needs of app firms. The broader problems of value creation will be solved. Right now, the bottlenecks in the system favor established firms over entrepreneurs, so there is an immediate advantage to value creation from established firms. That, too, could easily change through new innovation. Study of that awaits occurrence.

This gap between technical opportunity and market value creation is characteristic of information technology innovation over the last 60 years. Another element of continuity from the past is that the early uses of an important platform initiative need not be the ultimately most valuable ones, and that important interim innovation, even after early success, is important. The 21st century has brought a series of important changes to this. Some arise from scale, and the sheer size of the opportunity has drawn remarkable resources to the mobile area. Others arise from the quick entry of a second platform in competition with iOS/ITunes store, so that there is technical and market heterogeneity in even the general purpose components at a fairly late stage.