

This PDF is a selection from a published volume from the National Bureau of Economic Research

Volume Title: The Changing Frontier: Rethinking Science and Innovation Policy

Volume Author/Editor: Adam B. Jaffe and Benjamin F. Jones, editors

Volume Publisher: University of Chicago Press

Volume ISBNs: 0-226-28672-X, 978-0-226-28672-3

Volume URL: <http://www.nber.org/books/jaff13-1>

Conference Date: August 2-3, 2013

Publication Date: July 2015

Chapter Title: Economic Value Creation in Mobile Applications

Chapter Author(s): Timothy F. Bresnahan, Jason P. Davis, Pai-Ling Yin

Chapter URL: <http://www.nber.org/chapters/c13044>

Chapter pages in book: (p. 233 – 286)

Economic Value Creation in Mobile Applications

Timothy F. Bresnahan, Jason P. Davis, and Pai-Ling Yin

8.1 Introduction

No discussion of the great changes in the innovation processes of the twenty-first century would be complete without an examination of one of the newest growth poles: mobile applications. Mobile applications are software programs that run on a new class of mobile devices, smartphones, and tablets, which are typically connected to cell phone networks. The rapid growth of mobile devices has been accompanied by an equally rapid growth in app development, in substantial part because platform providers Apple and Google have lowered the costs of development and distribution of mobile applications.

Timothy F. Bresnahan is the Landau Professor in Technology and the Economy and a professor of economics at Stanford University, and a member of the board of directors of the National Bureau of Economic Research. Jason P. Davis is associate professor of entrepreneurship and family enterprise at INSEAD. Pai-Ling Yin is a social science research scholar at the Stanford Institute for Economic Policy Research.

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 data sets, 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 Kjartansson, Xing Li, Derek Lief, Sean Mandell, Laura Miron, Jaron Moore, Yulia Muzyrya, Abhishek Nagaraj, Joe Orsini, Francis Plaza, 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. We appreciate the valuable comments of Josh Lerner, Xibao Li, Ben Jones, Scott Stern, and Adam Jaffe. Pai-Ling Yin and Jason Davis benefited from the Karl Chang (1965) Innovation Fund and Edward B. Roberts (1957) Fund. This chapter represents data and conclusions as of its submission to the volume editors on October 13, 2013. For acknowledgments, sources of research support, and disclosure of the authors' material financial relationships, if any, please see <http://www.nber.org/chapters/c13044.ack>.

Like any new industry with significant promise, mobile apps have also engendered a long list of conjectures about where economic value might lie. Today, the industry is in the experimental phase of its life cycle. Like any information and communications technology (ICT) industry, much of the uncertainty and experimentation is about commercial rather than technical innovation. What kinds of apps will consumers use? Will consumers pay for them, or will they be advertising supported? Which existing industries will they “disrupt”?

Our goal in this chapter is to examine the supply of apps. We take up three topics that are suggested by the industry structure and by the early conjectures by industry participants:

1. Platform innovation: How is complementary innovation coordinated to create a new industry? How are scientific and technical opportunity linked to demand?
2. Industry evolution: How will experimentation in technologies, markets (and other institutions), and commercialization lead to changes over time?
3. Value creation: What could the largest new industry of our century contribute to economic growth?

Our investigation of platform innovation is laid out in section 8.2. Any innovation platform involves the sharing of general purpose components across applications; this sharing lowers entry costs for diverse, innovative applications. Our setting exhibits the most users of any application platform, the most apps, and the fastest growth thus far in ICT settings. The positive feedback loop appears to be working.

We choose to focus our study on the app developers and examine the rest of the industry participants from their perspective. A robust literature explores the technological history of mobile communication devices, platforms, and users, but the app developers (responsible for private and social value creation) are largely unstudied. Our empirical work has led us to a topic that cuts through all three of these areas: identification of the important bottlenecks to entrepreneurship, experimentation, and value creation that have arisen despite the considerable growth opportunity. The explosion in entrepreneurship has created an explosion in competition and an overwhelming choice set for the consumer. The problem arises in the first market institution required of a new mass market platform industry: the market that matches users and apps. We describe these institutions in section 8.3 and examine their implications throughout. We describe new data sets used in our analysis, some of which we have constructed, in section 8.4.

An application platform does not directly create economic value; instead, it lowers the costs of applications, enabling applications in a wide variety of sectors. We discuss this wide variety of applications sectors for mobile apps in section 8.5. Even though almost all of the most popular apps are not in market competition with one another, the size distribution of app demand is highly concentrated. This is true whether we look at downloads or usage. While in

principle this might arise because only a few highly popular app categories or apps have been discovered, we show in section 8.6 that there is a high rate of turnover (churn) in app success. We interpret the combination of concentration and churn as reflecting not only the underlying distribution of app attractiveness, but also competition across all apps categories for consumers' attention.

In section 8.7, we consider the implications of the "top lists" for industry development. We document a new, growing, and important category of apps, "corporate" apps, which form part of a consumer product and service firm's offerings to its customers. We show that these new, "nondisruptive" apps are considerably advantaged over entrepreneurial apps given the current state of the industry's market institutions.

In addition to experiments with a large variety of zones of application, developers are also experimenting with a wide variety of "monetization" strategies for apps, including paid apps, advertising-supported apps, "freemium" apps, and several more. We have gathered unique data on developers' commercialization strategies. These are analyzed in section 8.8.

To access most of their customers for a mass market mobile app, developers must write for the two largest platforms. A substantial body of discussion in the industry suggests that developers should write for iOS (Apple) first and then Android (Google) second. In section 8.9, we examine developers' platform choice behavior in two senses: platform preference (iOS vs. Android) and multihoming (both). Finally, we then consider alternative equilibrium scenarios for the industry in section 8.10.

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.¹ 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. Uncertainty about the value proposition for a new technology leads to exploration—not only at the scientific stage, but at the commercialization stage.² This is typical of general purpose

1. See Bresnahan and Greenstein (1999) for computing applications and Bar (2001) for communications industries. The widespread use of the Internet is taken up in Greenstein (2001).

2. The importance of a wide variety of *technologies* experimentally chosen by different firms early in the industry life cycle has been emphasized in a large literature, epitomized by Klepper (2002). We emphasize the parallel importance of a wide variety of commercialization experiments. In both cases, later market selection is critical. We also depart from the industry life cycle's standard modeling approach, which links the underlying uncertainty/variety/selection to horizontal industry structure variables, market shares, entry and exit, and so on. We emphasized different observables largely because we examine experimentation by firms that are not in direct competition with one another, so it is not clear what horizontal industry structure to examine.

technologies (GPTs), where the industry does not necessarily know all the uses or even the main uses at the beginning. 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

Our investigation of platform choice behavior by developers reveals a number of aspects of the barriers to application success in the current state of the industry. Entry barriers appear to be significantly lower for corporate apps than for apps from entrepreneurs. This arises not because of any discernible lack of technical or even marketing capabilities on the part of entrepreneurs, but rather because of the problems of matching users to apps. The current matching problem between developers and their customers slows the rate of innovation in mobile application, just as the low technical entry barriers raise it. The current matching problem also affects the direction of technical change, raising the fixed costs of entrepreneurial apps much more than those of corporate apps. “Disruption”—or even ordinary, high-value entrepreneurship—is disadvantaged versus continuity strategies of existing firms.

8.2 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, coordination issues arise among inventors of applications and/or of the general components.⁴ A new platform-based industry can sometimes recombine general purpose components that are already in existence. This, too, creates economies by avoiding reinvention 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, establish 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 technolo-

3. There is a large literature on this, summarized in Sichel (1997).

4. For the analysis of GPTs, see Bresnahan and Trajtenberg (1995) and Helpman and Trajtenberg (1998). The GPT literature emphasizes the market (uncoordinated) determination of the complementary rates of technical progress in a general purpose technology and applications. The two-sided market or platform-pricing literature (see Rysman 2009) emphasizes the use of a pricing mechanism, sometimes nonlinear, to coordinate supply and pricing and thereby internalize network externalities.

gies 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 as 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 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 the Android side.

The impact of all of this recombination and new invention was threefold. First, an application developer could create a new mobile app spending only a tiny fraction of the overall research and development (R&D) cost of providing it. Much of the R&D cost, including invention of mobile devices, mobile telephony transmission, the commercialized Internet, the cloud, and so forth, and the investment cost in infrastructure, including diffusion of mobile devices to users, had already been sunk and spread over thousands of applications. None of these common R&D or infrastructure investment costs were marginal to a particular app. For mobile applications developers, the fixed costs of offering a working system to users were dramatically lowered. Second, the potential economic return to new invention and investment in the platforms and in the preexisting complementary technologies that were recombined became higher. Third, applications developers were given free reign (and little guidance) to discover these opportunities. 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. In the case of mobile apps, this was achieved by the supply of a modest

5. Though there were technical antecedents to mobile applications (games and ringtones on Symbian phones, e-mail 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.

number of influential apps, which, taken together, were sufficient to attract a significant body of demanders.⁶ These included a media store, especially for music (iTunes), an app for accessing the Web (mobile browser), an interactive 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 killer apps has generated huge growth in mobile device demand. Today, more new smartphones (sixth year of diffusion) are sold worldwide than personal computers (PCs) (twenty-eighth year). For some kinds of consumers in some economies—for example, young adults in South Korea—the diffusion of smartphones has gone further than the diffusion of television. App developers have access to a very large body of demanders, as do advertisers seeking to run ads in apps.

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.

The rapid emergence of 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 8.1 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.

Android apps, measured by number, have caught up to iOS apps. One must be careful, however, about drawing any economic inference from a count of apps: the majority of apps are marginal and have not been downloaded or used by customers. We will revisit the question of the size of the supply of Android versus iOS apps later.

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. This enabled

6. This is not the only way to get over the hump into viability. A large demander, such as the 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,” that is, apps attractive enough to give users a motivation to buy the product.

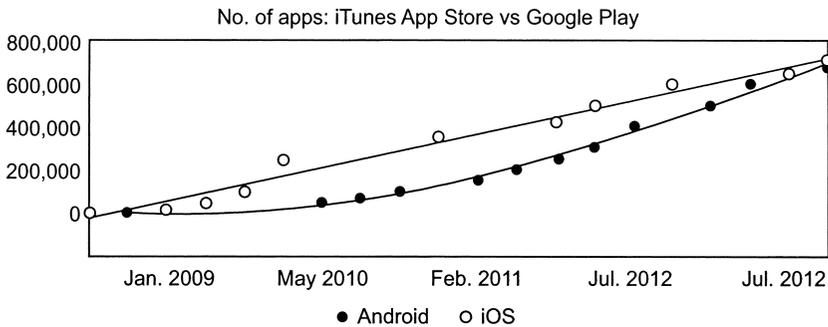


Fig. 8.1 Apps offered by platform over time

Source: InsideMobileApps.com.

(Accessed July 19, 2013: <http://www.insidemobileapps.com/2012/09/26/trackinggrowth-the-itunes-app-store-vs-google-play/>.)

an explosion of entry by app developers, a strong reflection of how low the incremental cost of app development has fallen.

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

8.3 Matching Apps to Customers: App Store Rankings

The mobile apps marketplace serves primarily consumers, and thus needs institutions to match applications products (and their sellers) to mass-market buyers. The design of a market institution for matching appears to be very difficult. New apps are introduced at a very high rate, and the problem of categorizing apps is largely unsolved. Moreover, the extremely rapid growth of the user base means that the distribution for demand for apps can be volatile. Within the broad established games category, products have short lives. Thus the search problem for a particular consumer is difficult. The solutions offered by the online stores, which we discuss in detail in this section, are of a “top list” form and do not appear to be very effective at matching demanders to desired apps.

8.3.1 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 vertically 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 distinct apps in July 2008; by May 2013 users had downloaded fifty billion apps.⁷ Apple's tight control over distribution gives it control over developer revenues, and Apple charges distribution fees of 30 percent 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, though it appears very difficult to categorize new apps. 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 trade-offs by displaying apps that have been most downloaded in the previous twenty-four hours on the "top" lists.⁸ 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.

An important implication of the top lists structure of the iTunes Store is that there is a strong winners-take-all flavor.⁹ An app can be popular because it is highly visible on the top lists, but it is on the top lists because it is very popular.

8.3.2 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 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 8.1).

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,

7. See <http://www.apple.com/pr/library/2013/05/16Apples-App-Store-Marks-Historic-50-Billionth-Download.html>. Accessed October 15, 2013.

8. At this writing (summer 2013), Apple was considering changing away from this ranking system to one also based on user reviews.

9. See Sorensen (2007) for analysis of this form of distribution mechanism.

10. Before March 2012, Google maintained separate app, music, and book markets. Afterward, the former Android Market merged with Google Music and Google eBookstore and became Google Play.

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 a user-action-based collaborative filter (recommending apps recently downloaded) and a user-opinion-based collaborative filter (comments and ratings).¹¹

Like the iTunes Store, Google Play has been very successful on a volume basis. But it is also problematic as a matching mechanism, for much the same reason. At this early stage of the industry, it is very hard to run an effective collaborative filter.

8.3.3 Problems Facing Collaborative Filters

Developers, especially developers who do not have a preexisting 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. 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 user. Apple is the only channel for distribution for its mobile devices, which exacerbates the problem of ranking being the dominant matching mechanism. Google’s openness and the existence of alternative app stores does not seem to be the solution to the ranking lists provided by iTunes. The matching problem has proved an as yet unsolvable challenge despite Google’s expertise in search. 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.

The problem of matching buyers and sellers is an industry-wide problem, not just a marketing problem for developers. Instead of competing with other apps in their same market, apps are competing with *unrelated apps* for consumer attention. Our empirical analysis will show that the store rankings fail to reflect the value created by the app and match heterogeneous consumers to apps. Although the entrepreneur could benefit from rapid adoption by the huge mass of end consumers, the ability to capture the “right” consumers in that dynamic is unclear, leading to lots of expensive

11. Apart from the star ratings and reviews, Google Play also has apps recommended by members of a user’s Google Plus “circles,” that is, their Google Plus social network. The importance of this improved collaborative filter is constrained by the limited penetration of Google Plus into widespread use as a social network.

investment in capturing all consumers. As a result, there is a curious duality about the new mobile platforms: they have lowered the *technical* costs of entering and supplying application products dramatically but have left very high *marketing* costs for entrant app developers. This duality is familiar from earlier ICT development platforms.¹² Further, it is familiar that marketing costs can affect industry structure.¹³ What is striking here is that the technical entry costs are so low because the platforms can offer so much accumulated and recombinable technology, while the marketing costs are so high, at least at this early stage, because the problem of matching app buyer and seller is particularly difficult when the seller is an end consumer.¹⁴

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. We will later discuss the emergence of a number of institutions as a result of the influence of collaborative filters. 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 platform providers are also tinkering with the structure of the online stores ranking systems to address the collaborative filter problems. For now, the rankings strongly influence the development of the app ecosystem.

8.4 Data

Our tables are built on three main data sources. The sampling frame for existing industry data sources are built around *products*, a peculiarity of our particular industry. There are no industry-wide data with the more usual economic sampling frames of firms or markets. We have thus worked to build firm-level and market-level data sets to pursue some parts of our analysis.

8.4.1 comScore

We utilize the “mobile metrix” data set from comScore. Like other comScore products, this is based on a panel of users, in this case approximately 10,000–12,000 US adult users of mobile devices. The bulk of the panel is

12. See, inter alia, Bresnahan and Greenstein (1999).

13. Fixed marketing costs have played a large role in the economic analysis of market structure and entry (see Sutton 1991). Sutton’s concept of “endogenous fixed costs” is closely related to the use of mass-media advertising by firms in competition with one another in product markets. We are examining a very different institutional structure in which products not in competition with one another in markets compete for the limited attention of consumers in a ranking system. These costs, too, are endogenous, in that the efforts of one firm to get attention for its products raises the costs of other firms’ getting attention for their products.

14. This is both like and unlike the problem of founding markets in earlier ICT technologies. There, the customer could typically be located, but the use of the technology involved invention of new organizations or work practices by the customer.

two subpanels, each approximately 5,000 users, one with Android phones and the other with iPhones.¹⁵

The underlying fundamental data are about each panelists' possession and usage of apps.¹⁶ However, the data come to us aggregated to the product*platform*month level.¹⁷

The sample of products (apps) on which data are available meet a minimum usage test for each 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 five (at least six) unique users.¹⁸ As a result, apps enter and leave the reported sample month to month.

It is very difficult to identify the same app on the two different platforms in this industry. In this data set, 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 that is different from its name in the iTunes Store and also different from its name in Google Play. An app that is only listed on one platform on comScore can also have a comScore name that is different from its name in that platform's app store. We have linked the comScore names of the apps to product and firm information.

For each month*platform*(included) app, comScore compiles a number of metrics. 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

15. There is also a smaller subpanel of iPad users, which we do not use in this chapter. The sample sizes change somewhat over the period in which the data exist.

16. There are also data on panelists' visits to mobile websites, but we do not use these data in this chapter.

17. 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 (a) they view the apps as different version of the same thing, (b) the app supplier views them as the same thing, or (c) 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.

18. 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 more 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.

Table 8.1 Number of unique apps by month

	Total	AOA	EOA	AOI	EOI	Multihomed
September 2012	1,301	715	439	862	586	276
October 2012	1,243	710	434	809	533	276
November 2012	1,202	691	408	794	511	283
December 2012	1,203	727	433	770	476	294
January 2013	1,231	793	507	724	438	286

Source: comScore.

Notes: AOA = all on Android, EOA = exclusive on Android, AOI = all on iTunes, and EOI = exclusive on iTunes. The table classifies apps in the comScore sample according to their availability on Android and iOS. 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.

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 United States 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 little more than 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 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 (0.06). 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 8.1 shows monthly statistics on app developers’ platform choices as reported by comScore between September 2012 and January 2013. Since the comScore panel, census projection methodology, and apps who use the comScore software development kit (SDK) are in flux over time in this very new industry, we are skeptical of reporting any changes 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

Table 8.2 Summary statistics of key variables

	Mean	Sd.	Min.	Max.
Unique monthly visitors (000)	1,289.92	4,324.60	11.12	60,805.94
Avg. unique daily visitors (000)	300.51	1,739.01	0.46	38,657.62
Total monthly minutes (millions)	1,993.24	16,289.40	0.26	391,802.25

Source: comScore.

Notes: There are 1,517 app-platform observations, January 2013. Unique monthly visitors provides a count of the number of comScore panel members who visit an app in the sample at least once per month, weighted to be representative of the US population. Average unique daily visitors provides the average of the daily unique visitors per day of an app over all days in a month. Total monthly minutes records the total minutes an app was used in month.

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, 1,301 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. “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.

In table 8.2, we report simple descriptive statistics for the size measures. One can immediately see that all variables are heavily skewed. For instance, on average an app is used for 1,993.24 million minutes per month. However, the standard deviation for minutes used is more than eight times larger than the mean. The maximum unique monthly visitors for an app is three to four orders of magnitude larger than the minimum.

8.4.2 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 chapter are download metrics for the online stores (iTunes and Google Play). App Annie copies the rankings of top apps from each of the online stores 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 twenty-four hours, whereas the top apps on Google Play are based on downloads in the past eight days.

8.4.3 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 are not collected anywhere else (they are not observable on any publicly available websites).

Our app questionnaire sample includes nearly 5,000 free apps (2,281 Android and 2,713 iOS). We have two sample inclusion criteria. First, we attempt to survey the most popular apps. This is simple 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 that has had at least 500,000 downloads. It is more difficult on the iTunes Store, which has no similar data publically visible. For iOS, we define a popular app as one that has been on the App Annie rankings list for at least ten days at any time. We continue to survey popular apps today, but the samples used in the tables in this chapter focus on three periods: September 2012–January 2013, January 2013, and June 2013.

We also include apps in the questionnaire if the app appears in the comScore data. This adds a surprising number of apps (about 1,000). This could arise because comScore is a measure of current usage while our definitions of “popular” are based on downloads, or because of other gaps between comScore's sampling frame and the marketplace. However, we think it is more likely that it reflects the natural sampling uncertainty. To be popular in our definitions includes a large number of quite unpopular apps, as we have pushed our definition of popular far out into the long tail of apps. For these observations, comScore's definition of “widely used” as 1/1,000 in a sample of 5,000 has considerable sampling variability.

For each app, the questionnaire asks whether the app utilizes advertising and a number of questions about the products and services that are 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 chapter.

We see an app on either iOS or Android. We ask the research assistant to look at the developer's website to see if the app is present on the other platform.¹⁹ While on the developer's website, we ask whether the developer has a different app on the other platform. Similarly, the website tells us whether the app is part of a “freemium” pair. Our examination of the developer websites is the closest thing to a firm-level data set on app suppliers in this industry.

19. Since developer websites were not always available, this part of the questionnaire data set is smaller than the rest.

There are two problems with our questionnaire sampling frame; neither is solvable. First, we do not know exactly what function of downloads we are using as a cutoff for iOS apps. Second, if advertisements are targeted, then our data on advertising is conditional on our survey respondents' behavior (i.e., students who download many apps per week) and thus may not be representative of the market. At this stage of industry development, our conversations with advertising firms suggest that targeting at the individual level is not commonly used, so this second sampling problem is less important.

8.5 App Success is Highly Concentrated

Many theoretical models of platform success simply count apps.²⁰ Mobile industry sources do as well, noting that there are approximately three-quarters 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.

Mobile apps compete in a wide variety of markets: games (including many subcategories), entertainment (music, television, books, or magazines in any of a wide variety of business models), content (news, weather, mapping services, financial information), Internet services (browsers or maps), travel services (airline and hotel online resellers like Expedia), online advice (like tripadvisor), communications services (text-message managers, Skype, social networks), and so on. This heterogeneity in purpose across the hundreds of thousands of apps has several implications for our analysis. From an industry life cycle perspective, the apps world does not call for one set of experiments to find the best app technically and in terms of business model, it calls for dozens of parallel sets of experiments. Second, most apps are not in *market* competition with a large number of other apps. In contrast, all apps are in competition with all other apps for users' *attention* in the app store.

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.²¹ As we noted in the data section, there is a very

20. Farrell and Klemperer (2006) give an overview of these models.

21. 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.

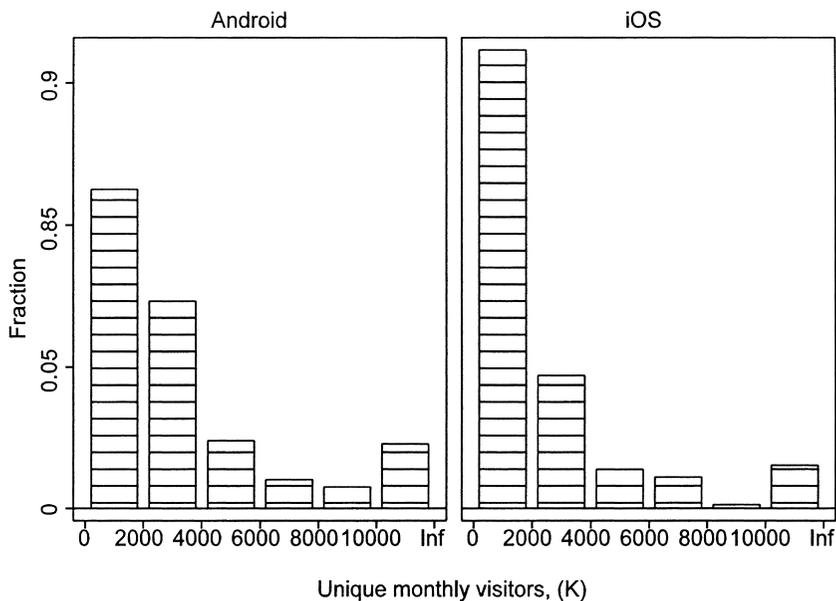


Fig. 8.2 Size distribution of apps by platform

Source: comScore, January 2013. This histogram of app usage has equally spaced bins, except for the rightmost one, which has no upper limit. The vertical axis is the log of the fraction in the bin.

substantial correlation among a number of these measures, and we are thus drawn to examine only two of them: (a) how many users have the app on their device and use it, and (b) how many minutes a typical user has the app open. We begin by looking at the first, measured by “total unique monthly visitors” for January 2013.

This investigation leads us to two facts, each about a different portion of the cross-section distribution of app demand. First, both platforms exhibit a high degree of skew in the size distribution of app demand. We noted above that apps compete in hundreds of markets. With only a couple of exceptions, the top twenty apps on each platform are not in competition with one another in those markets. We will return throughout this chapter to understanding the economic forces behind this high degree of concentration across all categories.

A similar phenomenon arises across the range of app demand high enough to be observed in comScore, that is, showing up on about 1/1,000 phones. To see this, we examine the histogram of the size distribution of apps (still measured by January 2013 “total unique monthly visitors,” but in log scale) shown in figure 8.2. As you can see, the frequency of each fixed-width-bin of projected total US users drops rapidly as you make the number of users larger. What does the histogram *not* show? Stronger versions of the same

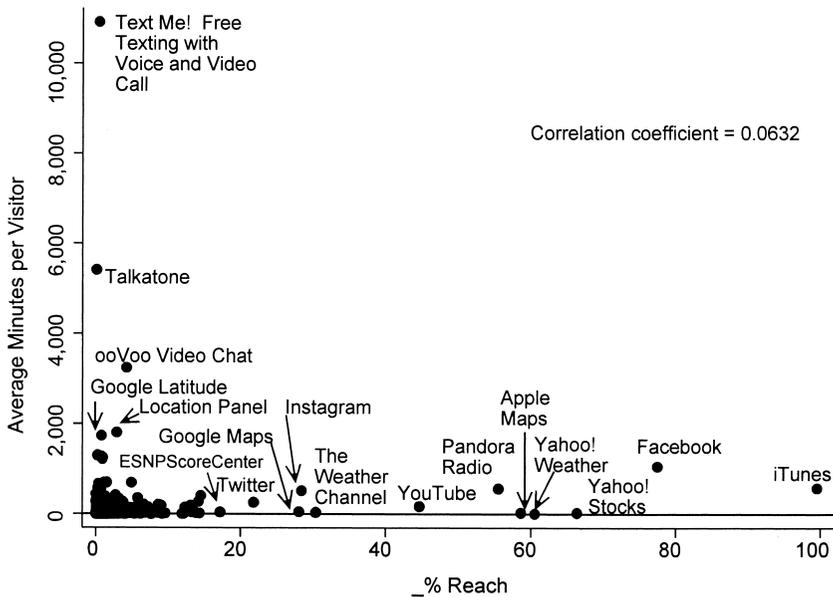


Fig. 8.3 Joint distribution of minutes and reach

Source: comScore, January 2013.

phenomenon. The rightmost cell contains the very few apps that have the most usage. To the left of the leftmost cell lie the overwhelming majority of apps (all but about 700 of the 750,000 or so on each platform) that have too little usage to appear in comScore.

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 8.3.

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).²² Industry participants, not surprisingly, think of this as a distinction between apps, which deeply engage the user (y axis) and apps that have mass appeal to a very large number of users (x axis). It is clear that these are not the same object. Nonetheless,

22. 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.

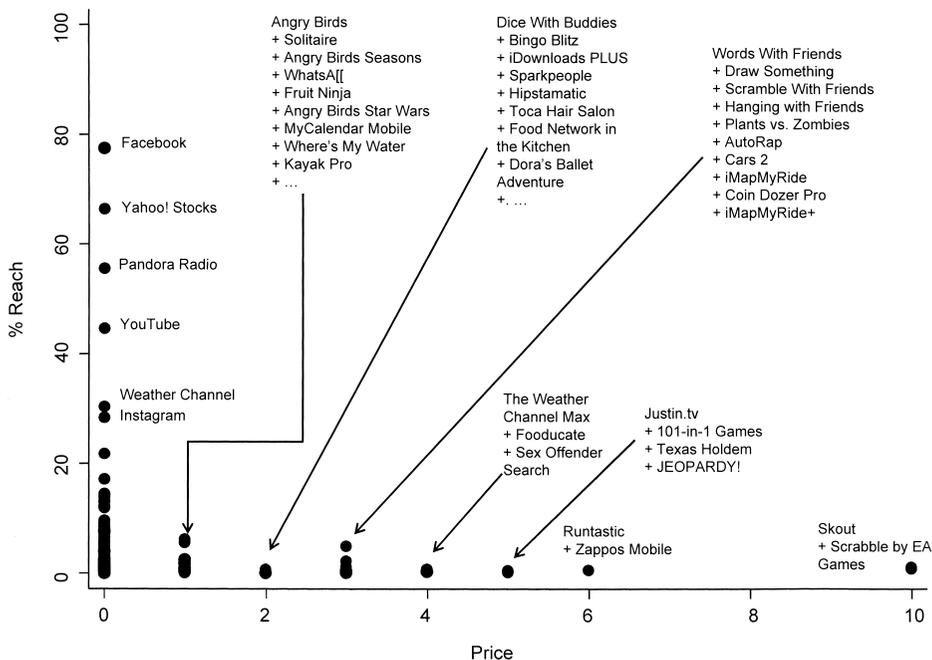


Fig. 8.4 Joint distribution of price and reach

Source: comScore, January 2013, and iTunes.

it is easy to see that there would be a very highly skewed distribution of the attractiveness of apps whether that was measured by reach, minutes, or both.

At least one source of variation in the size distribution, pricing, can be descriptively examined. In figure 8.4, we look at the joint distribution of reach and of price for the iOS platform in the same month. This is once again based on the comScore data, which we have linked to the iTunes Store in order to collect the prices.²³ 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, if they are high enough up in distribution to be labeled in the figure, are most frequently games.

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

23. 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.

are the online stores, and those add an element of positive feedback to app demand through the online ranking system. No reader of Sorensen (2007) 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. We will return to this third, institutional, cause over and over throughout this chapter.

8.6 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 a period of time. In this section, we examine that proposition through analysis of the churn in app rankings at different time scales and for both Android and iOS apps.

To examine churn, we adopt the simple short-run dynamic measure of Waldfogel (2013), that is, “the number of apps that appear on the top fifty list seen at two dates.” This is unlike the question of how many top fifty economics departments or business schools there are: at any given date there are only fifty apps on the top fifty list. However, if we look at two dates, there may have been some turnover. If, for example, (exactly) seven apps that were on the top fifty at time s are off it at time t , fifty-seven apps total will have been on the top list at the two times. The higher this number, the more churn there is.

We measure churn in the flow-of-apps-downloaded based on administrative records gathered from the iTunes Store and Google Play by the recommendation site App Annie.²⁴ App Annie 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 a more recent period in which both stores were open.

More precisely, using daily rankings from App Annie, we determine the volatility of a ranking from two dates s and t as follows: take one of the ranking methods from the online stores, “top paid,” “top free,” and so on. Locate the top fifty apps in both rankings, R_s and R_t . Next, we count the number of apps in the union of the two sets, D . Table 8.3 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.²⁵

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 Android “top fifty paid”

24. 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 for existing users is very high, so the stock does not lag all that far behind the flow.

25. 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 (IAP), 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. Since there are no ties in the data, the reported means all lie between fifty and one hundred.

Table 8.3 Average number of distinct apps in top fifty daily lists at different lags

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

Source: App Annie rankings between January 1, 2012, and June 28, 2013.

Notes: This table shows, for a number of different daily top fifty rankings, the average number of apps that make up the rankings on two different days. The columns index the number of days between the rankings. The top-left figure shows that the Android top fifty paid ranking on two consecutive days contains on average 51.35 apps.

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 fifty list.

The first fact visible in table 8.3 is that there is a good deal of churn in all the lists. This is perhaps what you would expect from a very new market. A second fact can be learned by looking at the churn that happens within a month (third column) versus the churn that happens within a year (last column) for all the lists except the Android “new” lists.²⁶ The churn is in general high frequency, with quite a bit of the turnover that happens in a year occurring within a month.

A third fact is that after either a day or a week the churn 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 is surprising: new apps are experiments, and they mostly fail. Many developers tell us that the Android top free lists are based on an eight-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.

Fourth, 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 as we increase $t-s$. After a year, for example, just over 50 percent of the top free Android apps (78.15–50) will have departed, while over 80 percent of the top free iOS apps (91.24–50) will have departed the top lists. There is an obvious explanation for this. 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.

26. By construction, all the turnover on the Android “new” lists happens within a month.

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

There are, of course, a number of possible explanations for these simple descriptive statistics about churn. New and better apps could be constantly introduced, displacing old (ninety 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. These explanations are unpromising. Developers and other market participants have not proposed these explanations, 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 fifty 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.

8.7 “Top List” Implications for Market Development

We have established that the collaborative filter of the online apps stores is an important barrier facing apps. Many entrepreneurs entered this business hoping to take advantage of the technical opportunity to attract consumers and then sell advertisements to firms who were hoping to acquire those consumers as customers. Difficult search and matching makes the customer acquisition part of commercialization difficult for an entrepreneur who has written an attractive app. If consumers who would value the app cannot find it, demand will be small. This gives app developers an incentive to build volume quickly. The high marketing costs have led to a market response, with a host of services that have arisen to try and solve commercialization and monetization problems for apps struggling to become visible in the clutter of apps.

This can lead to the irony, (now) frequently noted by industry observers, of an app that plans to be supported by advertising in the future but spends on mobile advertising today to get more downloads despite zero revenue. 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 another’s 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. Regardless of the channel, developers have allocated resources to purchase advertising to promote apps instead of attracting marketing dollars from products outside of the industry.

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 in-app purchasing (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 that 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 that 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 of using apps and spending money inside apps.

The high marketing costs of using the app stores fall more severely on entrepreneurial app developers than on existing firms with consumer connections. The costs of inducing many consumers to download an app are much smaller for those app developers who already have a marketing connection to their customers, typically because those are already *online* or *product* or *service market* customers. If this is correct, it suggests a market equilibrium shift away from “disruptive” entrepreneurs to existing firms.

8.7.1 Evidence of Buying Users

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.²⁷ We count all the different ads that show up in one use of the app—defined as, at a minimum, passing all the

27. 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).

Table 8.4 **Banner ad content**

	Android	iOS
Apps with banner ads	$N = 1,106$	$N = 1,583$
Any app ad	.802	.847
Any non-app ad	.370	.371

Source: App questionnaire as of January 2013.

Note: Column totals sum to more than one because an observation is an app and the measures are for *any* advertisement of a given type, so that multi-ad apps may show up in both rows.

Table 8.5 **Percent of apps with banner ads in each category**

	App ad	Non-app ad
Same firm	0.348	0.091
Different firm	0.623	0.295

Source: App questionnaire as of January 2013.

Note: Column totals sum to more than one because an observation is an app and the measures are for any advertisement of a given type, so that multi-ad apps may show up in both rows. Figures are weighted average over platforms.

places where ads might be shown. Thus the probabilities sum to (considerably) more than one. The 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.

A surprising finding from table 8.4 is that if an app has a banner ad, there is an 80 percent chance or greater that it will serve an ad for an app! In table 8.5, 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 percent. The demand for ad space in apps is demand for space to advertise other mobile apps.²⁸

In table 8.6, 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 others. We recorded data for more categories than this, but once the table is viewed it will be obvious why this is what we report.

This is a complex enough table that it is useful to walk through the numbers. Table 8.6 reports the products and advertisers seen in Android and iOS

28. 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.

Table 8.6 Advertised products and vertical integration in advertising

	Android banner ad apps <i>n</i> = 1,106		iOS banner ad apps <i>n</i> = 1,153	
	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 as of January 2013.

platform apps that displayed at least one banner ad. Within the Android banner apps, about 10 percent (.102) display a banner ad for another app from the same firm, while just about 2 percent (.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 8.6, 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 50 percent (.486) display a banner ad for another app from the same firm, while over 10 percent (.133) display a banner ad for another product from the same firm that is not an app. Only 57 percent 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 (a) 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 twenty-four hours rather than the last week of downloads, and (b) since iTunes is the only distribution channel for iOS apps, influencing iTunes rankings is relatively more important to an app's 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, that is, 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. See figure 8.14.

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 does not rebut this finding. Those numbers come from summing the revenues across app developers,²⁹ 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.

8.8 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.

Within all of these three broad categories, early app development was heavily influenced by the idea that app developers would be *entrepreneurs* and that

29. “Revenues” in the industry sources are themselves not the strongest numbers in the economy.

they would seek to *disrupt* existing industries.³⁰ The scope of the anticipated disruption was dramatic. Entertainment, advertising, and media markets would be disrupted. So, too, would industries that could take advantage of a very localized advertising and information service; in retail trade, for example, attracting consumers into the store might depend less on the brand capital of a chain store or on the locational attractiveness of a mall, and more on sending an ad to someone walking down the sidewalk. In parallel, industries characterized by transaction costs of aligning buyers and sellers—say taxicabs—might be disrupted by new services to match buyers and sellers over mobile devices.

There are at least two reasons to anticipate some disruption. First, earlier advertising and communications technologies have changed industry structure in advertising-intensive industries and in communications-intensive industries.³¹ Second, the creation of the mobile app development platforms has certainly disrupted mobile telephony, as influence over customers and over the direction of industry standard setting has shifted from the mobile carriers to Google and Apple. One can only hope for more disruption as mobile devices compete with established forms of mass market computing such as PCs. It is worth noting, however, that essentially none of this earlier disruption came from entrepreneurs.

The broadly anticipated categories of app value creation also suggested some variety in how apps would earn their economic return. Early conjectures focused on two main ideas: consumers might pay to be entertained or to avoid the hassle of transaction costs, or advertisers might pay to gain consumers' attention.

Given the heterogeneity in apps, it is not surprising that there is a good deal of heterogeneity in the way the developer seeks to earn an economic return if the app is successful. It should be noted that although some monetization strategies like freemium have been featured prominently in the press, many entrepreneurs are uncertain about which methods are most effective for their apps. For example, one iPhone game developer echoed the sentiment of many entrepreneurs we spoke with:

In all honesty, we still haven't figured out whether ads, freemium, discounts, or in-app purchases are the best way to make money off our loyal customer base, so we've committed to a period of experimentation with these different methods. It may come down to how these strategies work in combination with our different games. Maybe we can have a menu of payment options, but I suspect we'll have to commit to one or two because our customers may revolt if it gets too complicated.

As may be apparent, this is already significantly broader than that suggested by the “who pays” dichotomy between advertising-supported and

30. See, for example, Christensen, Johnson, and Horn (2010).

31. See, for example, Bresnahan (1992) and Sutton (1991) for the impact of mass communications on industry structure.

paid-by-the user apps. We now review a number of “monetization” strategies taken up by developers. Once we have discussed the main forms of app monetization, we will present some simple statistics on their prevalence.

8.8.1 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 users’ lives. For a number of reasons, however, straight-up paid apps are not the most common revenue-gaining strategy, even when the user ultimately pays.

Two different strategies for charging the user have a try-before-you buy structure. “Freemium” apps have two versions. The “lite” version is free but limited in some way. A “pro” version is available at a positive price and is typically recommended with the “lite” version and sometimes even purchased within it.

Similarly, in-app purchasing (IAP) lets users buy complements to the app as they are using it; in a war game, IAP could be used to buy a bigger sword or more lives for the player’s avatar, for example. The IAP 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.³²

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 up front. Furthermore, apps appear to be experience goods, as evidenced by the large number of apps in our comScore data that get downloaded but are not used.

Other apps use a subscription model, a recurring monthly or yearly IAP 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, antivirus, and even navigation apps have also called for users to subscribe.³³

32. It is not uncommon for apps that monetize through IAP 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 IAP makes acquiring the scarce currency much easier. Further investigation could help us better understand the in-game currency dynamics, especially as they relate to dual-currency games.

33. 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 preestablished username/password for authentication. The carriers have even sponsored some subscription apps that allow users to pay 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.

8.8.2 Advertising

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, and so forth—and the broker sells that space to potential advertisers. The app developer can be paid either like traditional media firms, that is, by the number of users who are shown the ad or like traditional online firms, 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.

The wide variety of organizational structures for advertising is accompanied by 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 (ads that 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 entire 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 of the still-visible app.³⁴

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 touch in the course of playing 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 location, past purchases, the content of communications, and so forth. 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, and so forth 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

34. 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.

typically quite primitive. While some advertising is integrated with the users' past behavior, location, and the like, 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.³⁵ Ad networks, including those that are part of the large platform firms, are going to need to work out what to measure about users, how to sell that knowledge, and how to divide the rents associated with targeted advertising with developers. Users will need to decide whether to respond to 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) 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. 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 percent share brand that emerged out of a very fragmented US beer industry). For mobile advertising, those coordinated commercial inventions (as opposed to technical inventions) still lie largely in the future.

8.8.3 No (Current) Revenue Stream—Corporate Apps

Finally, there is a very important class of app that 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 his or her flight.

These corporate apps are large and growing as a portion of all app downloads. By the time this chapter is published, this 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

35. A number of brokers and advertisers have told us that making the target audience part of app design (to the extent an advertiser cares about the message recipient) lies largely in the future.

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 chapter 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.

8.8.4 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 venture capital 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 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 advertising-supported apps may seek to avoid annoying users in early stages, that is, 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 initial public offering (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 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. We have 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. Their app solved a problem for them, and they are quite pleased that it solves a similar problem for others. These developers may be anticapitalist 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. Our industry sources tell us to expect such “hobbyists” more (today) on Android than iOS.

8.8.5 Some Statistics

In this section, we quickly examine some statistics about app developer monetization strategies. We first look at figure 8.5 and figure 8.6, which come from an industry source.

Note that the category of corporate apps is not present, as the Developer Economics data source follows current industry practice of focusing on “monetization” as a topic in management-normative analysis. The survey shows that there is considerable heterogeneity in the monetization strategies used by developers. While advertising is the most popular, there is also considerable use by developers of paid apps, in-app purchases, and/or free-mium. Subscription payments used significantly less.

In contrast, when we look at the revenue-per-app figures, subscription apps are the highest, with in-app payments second. Advertising-supported apps have the lowest revenue. The management-normative conclusion that developers should shift out of advertising-supported apps and into subscription apps is likely not warranted. Subscription models, as we pointed out above, are only suitable for certain kinds of apps, most obviously those that distribute content such as music.

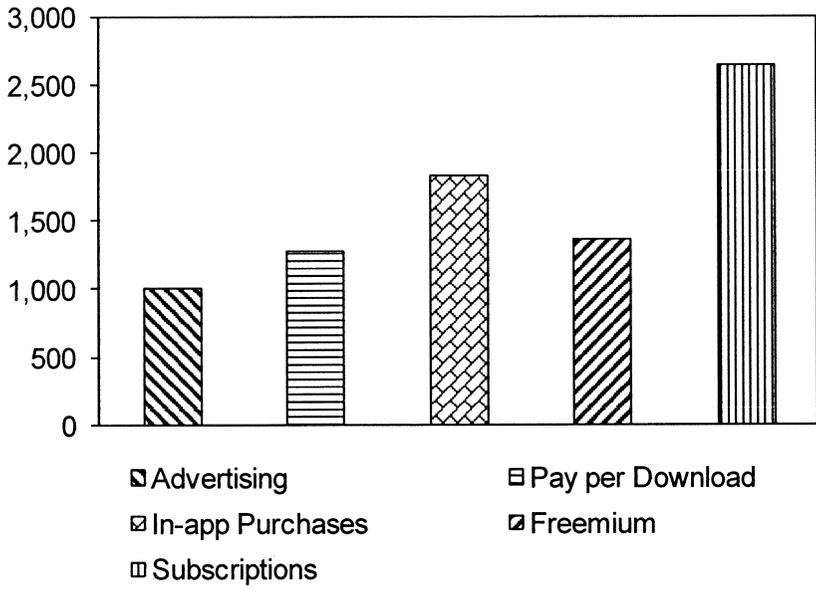


Fig. 8.5 Revenue per app by monetization strategy

Source: Developer Economics survey.

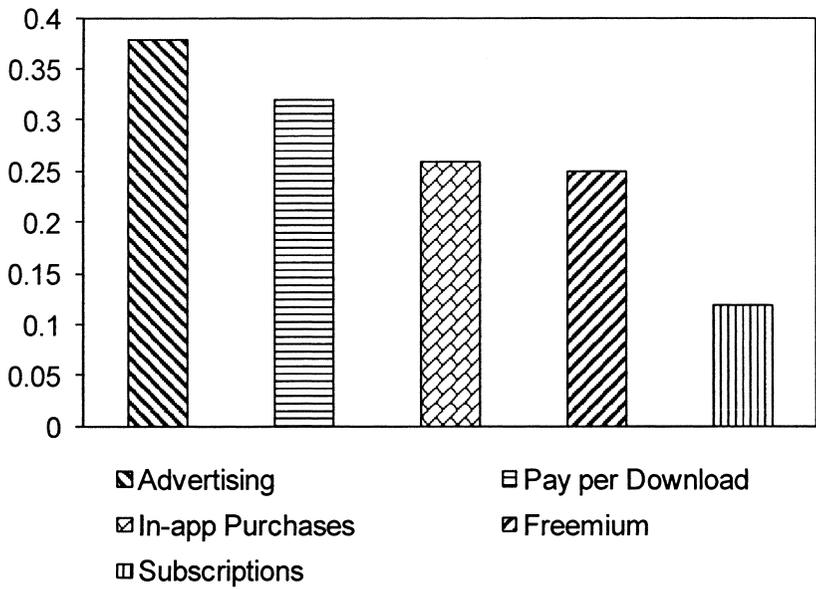


Fig. 8.6 Frequency of use of monetization methods

Source: Developer Economics survey.

Table 8.7 Frequency of use of IAP and advertising in free apps (wide sample)

	(1)	(2)	(3)	(4)
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 (July 2013) sample is all apps. N: Android 2,281, iOS 2,713.

On the other hand, these figures do suggest that the question of who pays for mobile apps—users or advertisers—is not yet settled. Since there is some overlap between categories (e.g., an app can have both advertising and IAP), it is a slightly odd calculation to look at the share in total revenue represented by a category. Still, it appears that the total revenue across apps that have advertising is smaller than the total revenue of apps that use any other method. Advertising-supported apps are not taking over, at least not in a revenue sense. Given that a good portion of the advertising revenue of the typical app comes from ads for other apps, the overall net advertising revenue of the app industry must be even smaller.

We can get another estimate of the relative size of the different categories of revenue strategies, and include the important “corporate” form, by examining our app questionnaire data. In table 8.7, we show revenue strategies for the most popular free apps on the Android platform, and for a sample of iOS apps that 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: (a) neither advertising nor in-app purchases; (b) no advertising, but in-app purchases; (c) some advertising, but no in-app purchases; and (d) some advertising and in-app purchases.

On both platforms, the most widely used monetization strategy for free apps is advertising without IAP, which was used in 2012 for 44.9 percent of free Android apps and 39.4 percent 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 percent on iOS compared to 7.98 percent on Android). The higher use of IAP on iOS is consistent with industry surveys 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 percent) of free Android apps neither use advertising nor in-app purchases to users, and similarly for iOS, about a fifth (22.1 percent.) This highlights the significance of the class of apps that do not yet pursue a stand-alone commer-

Table 8.8 Frequency of use of IAP and advertising in free apps (narrow sample)

	(1)	(2)	(3)	(4)
Ads used?	No	No	Yes	Yes
IAP used?	No	Yes	No	Yes
	Neither	IAP only	Ads only	Both
Android and iOS	.258	.135	.357	.251

Source: App questionnaire (July 2013) sample is apps in comScore.

cialization strategy that is tied to the app economy.³⁶ These are mostly corporate apps.

As a robustness check, table 8.8 presents the same set of frequencies as table 8.7, 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 8.8 are the apps that satisfy the intersection of the two criteria: both these apps are top apps and appear in the comScore data. Furthermore, the subset is limited to those apps in this intersection that 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 8.7.

We have also checked the “neither” column against a measure of firm type. Overwhelmingly, these apps are associated with a firm that, in addition to having the mobile app, is a nonmobile consumer products and services (pre-existing, “corporate”) firm. At this stage this is a limited subsample, since we are in process of defining firm types after linking each app to its sponsoring firm.

The emergence of a large category of “corporate” apps was a surprise. Most industry figures expected a primarily entrepreneurial form of app supplier. Further, the most common articulation of the broad value proposition of mobile was that an app developer’s knowledge of where the consumer was located would permit selling location-specific advertisements. The mobile opportunity, in this advertising-centric view, was a more granular version of zip codes. However, mobile devices and services are a general purpose technology. Consumer product and services corporations have adopted this GPT and used it to provide customer service. This is an important, unanticipated growth category and represents a potentially large-scale increase in the economic value of the mobile opportunity.

36. 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.

Table 8.9 Asymmetries between platforms

iOS and iTunes Store	Android and 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:	Changes in environment from year to year
Porting an app to the newest iPhone/iPad devices from older ones usually simple	(e.g., substantial UI changes) Different hardware manufacturers use different OS versions
Distribution restricted to iTunes Store	Open; multiple distribution channels

Source: Discussion with industry participants.

8.9 Developer Behavior: Platform Choice and Multihoming

In this section, we examine the platform choices of smartphone app developers on iTunes and Google Play. We take up three topics. First, we examine the relative attractiveness of the two main platforms. While much industry discussion suggests “iOS-first” as a strategy, developer platform choice behavior is close to evenly balanced between iOS and Android at this time. Since the installed base of iOS and Android phones is approximately equal, this is consistent with the theory that the most important determinant of developer behavior is the installed base. Second, we examine developer decisions to multihome, that is, write for both platforms. Here the basic facts are more complex. If we look at the decision to *write* the app for both iOS and Android, multihoming is quite common. However, if we look at *marketing* the app to a mass market, multihoming is rarer, about one-third as common. Our interpretation of this dual result is consistent with the theme of this chapter: at this stage of industry development the technical entry barriers to either platform are low for developers, but marketing barriers remain high because of the problem of matching app demand and supply. Third, we consider the possibility of future changes in the platform market.

8.9.1 Relative Attractiveness of the Two Platforms

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 8.9.

The industry discussion summarized in this table does reflect elements of the usual economic theories, that is, the size of the installed base of devices on each of the competing platforms. At present, an app developer can reach approximately half of all US smartphone users by writing only for iOS and distributing through the iTunes Store, and approximately half of all mobile users by writing only for Android and distributing through the Google Play store, Amazon, and so on. A much smaller installed base of users can be reached by writing for some other platforms, such as Windows mobile, Blackberry, and so on. Thus, one attractive theory is that developers ought to be approximately indifferent between the two platforms.

The dynamics of platform competition might lead to different expectations by developers of future installed base sizes. Apple's iOS started with a large lead, but Apple left an opening by not having iOS devices universally available. Android phones then grew more rapidly. Today, there is some discussion of possible future Android growth, but hardly a wide consensus on the part of the developers that either platform will be important in the future. On the other hand, developers do appear not to expect that a third platform will catch up to today's top two, though both Microsoft and Blackberry have proposed one.

There is also a great deal of discussion of the per-customer profitability for a developer on each platform. A broad consensus in the industry appears to favor iOS over Android here, primarily because iPhone users spend more through their mobile devices. It is thus less than obvious whether the initial lead for iOS plus the richer potential customers or the open systems approach of Android will have the advantage.

Examining developer platform preference by examining developer platform choices is, in general, a difficult problem.³⁷ We undertake a simple version based on the comScore sampling frame. In comScore, an app appears if 1/1,000 of panelists use it. The iPhone panel and Android phone panel are the same size. One simple approach to examining developer choice is to count the number of apps that appear in the iPhone and Android phone panels. This is shown in figure 8.7, which shows a very interesting trend. In September 2012, at the beginning of the time period shown, which is also likely the earliest date at which the comScore data can be used for this purpose, there appears to be more developer preference for iPhone apps than for Android apps. By January 2013, the most recent data that we have in cleaned form, this preference appears to have been reversed. This pattern is, of course, consistent with the idea that Android is catching up or even overtaking iOS in attractiveness with app developers, as it clearly is in attractiveness to users. The small differences between the two platforms are also consistent with the theoretical prediction that, with approximately equally

37. Rysman (2009) surveys the literature very well.

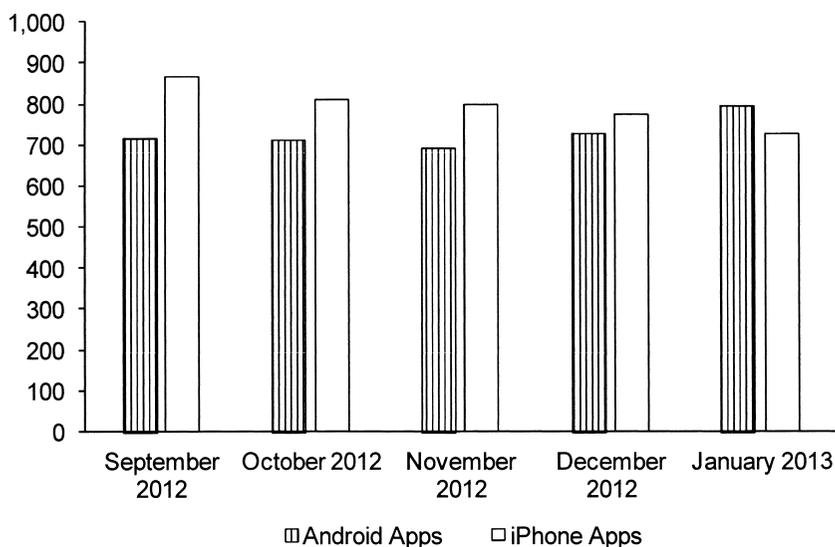


Fig. 8.7 Number of apps by platform by month

Source: comScore. Data are a subset of those in table 8.1.

many users, the two leading smartphone platforms should be approximately equally attractive to developers.

As of this writing, it appears that Apple's original lead has disappeared. Whether the mild trend to Google visible in the total number of apps (figure 8.1) or in the number of apps with reach over .001 (figure 8.7) is the start of a platform shift depends on the attractiveness of the trending apps, a topic for the future.

8.9.2 Multihoming

Multihoming—supplying the same app for more than one platform—is another aspect of developer supply behavior with important implications in the economics of platform markets.³⁸ With two platforms of approximately equal installed base, multihoming is also important for commercialization. Any app developer attempting to supply a mass market must multihome to both iOS and Android (Blackberry and Windows mobile do not contribute much installed base at this juncture). It is clear that whether, when, and how to multihome are critical decisions facing new ventures and established firms producing mobile applications. For example, one developer of productivity applications summarized the trade-offs facing his firm during multihoming:

38. There is a large literature about competition for developers and completion for users by platform firms. One behavior emphasized in this literature is multihoming. See, for example, Armstrong and Wright (2007).

Like many companies, we started on the iPhone, but with the explosion of the Android market, we are seriously considering a move. We don't have the infinite resources of a Zynga or Rovio, so the real decision is whether to make a new iPhone app for our existing customer base, or get in on the ground floor with Google Play by porting our existing app to Android. The problem is: there is a lot more uncertainty about whether Android will ever make money, but it's possible that iTunes is tapped out for us.

Analysis of multihoming also calls for matching apps across the two platforms and across data sources. Matching apps across the two platforms is significantly more difficult than one 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.³⁹ For the limited set of apps it covers, comScore has linked across platforms, but unfortunately comScore does not use the same firm and product names as do developers. Our best results have come from looking at developer websites one by one and capturing all the apps—and their unique identifiers on each platform—listed by the developer.⁴⁰ Since this website is found through our app questionnaire, this hand-match lets us employ our app questionnaire and the comScore data to examine the current state of multihoming.

Finally, we define two different definitions of multihoming and use two different data sets to examine them. One definition of multihoming is “written for both platforms.” This technical definition of multihoming is observed in our developer website survey. Another definition of multihoming is “marketed successfully on both platforms.” We introduce this second definition because we suspect that the fixed costs of marketing a mass-market app may, along with the fixed costs of writing it, limit entry. This technical + marketing definition of multihoming is observed in the comScore data. Finally, we should point out that all our data sets condition on app success; in what follows, we will be looking at multihoming among apps that have had considerable success on at least one platform.

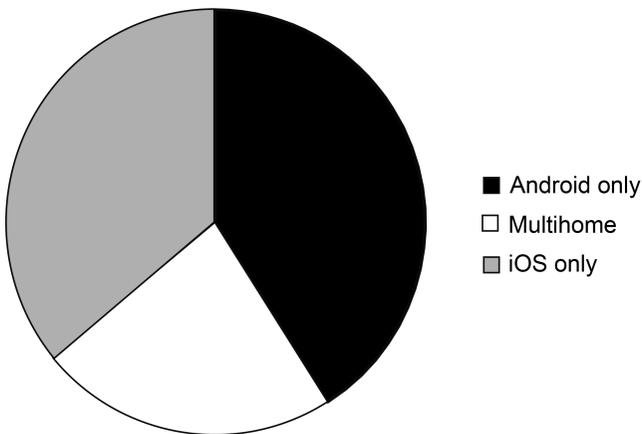
We begin with the January 2013 comScore sample.⁴¹ In that month, there were 1,231 apps in the comScore data set. In figure 8.8, we show the distri-

39. 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 half of true matches—that is, have more false negatives than true positives—and also identify a significant number of false positives, about one-tenth of the number of true positives.

40. 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 handmatched comScore “apps” to unique identifier app ids. We very much appreciate the help comScore has given us in resolving the last few difficult cases.

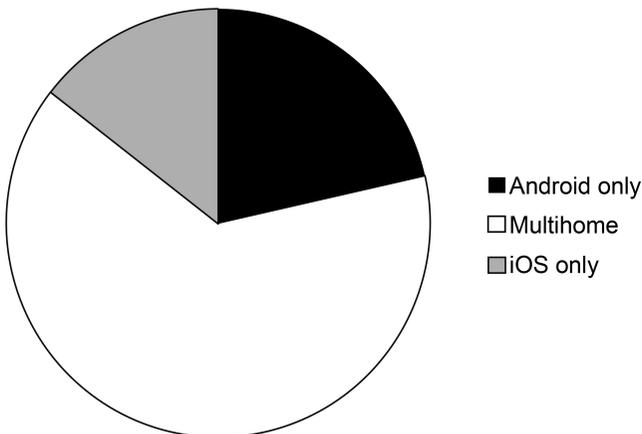
41. The comScore mobile product is new and had an initial period of rapid improvements in coverage; January 2013 is after the comScore mobile sample and definitions settled down.

Platform Choices in Marketing Sense

**Fig. 8.8 Multihoming, iOS exclusive, and Android exclusive apps in comScore**

Source: comScore, January 2013. "Multihoming" here means appearing in both comScore panels.

Platform Choices in Technical Sense

**Fig. 8.9 Multihoming, iOS exclusive, and Android exclusive apps in comScore**

Source: App questionnaires (platform choice based on listing of app on developer website).

bution of apps among iOS only, Android only, and both categories. As you can see, there is a low rate of multihoming.

We can take a broader view by drawing on our questionnaires. Figure 8.9 is based on the same sample of 1,231 apps as figure 8.8. Here, however, we define multihoming based on the developer-website information. An app is multihomed if the developer has links to download it both from the iTunes Store and Google Play. A single-homed app is one that appears on comScore

on one platform, but the developer does not list a link to the other platform on its website. As you can see, the rate of multihoming on this “technical” definition is dramatically higher.

The big difference here is between the tendency to *supply* an app for both platforms, which yields a rate of 64 percent multihoming (figure 8.9) in our survey of developer websites, and the tendency to *widely market* the app on both platforms, which yields a 23 percent multihoming rate in the comScore data (figure 8.8). We interpret this large difference as 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.⁴²

8.9.3 Weighted Multihoming Rates

We noted previously that the size distribution of app demand was highly skewed. The larger the demand for an app, the more likely it is to be profitable for the developer to bear the incremental fixed costs of multihoming, that is, the technical costs of porting to the second platform and the marketing costs of establishing a connection to consumers on the second platform.⁴³ 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. We examine this hypotheses once again in the January 2013 comScore data.

To associate weights with each app, we now turn to a frame in which each app has a “base” platform and the definition of multihoming is that the app is also observed on the other platform. In figure 8.10 the base platform is iOS, so “multihoming” means availability for Android. Figure 8.11 is symmetric; the base platform is Android and multihoming means we observe the app on iOS.

In both figure 8.10 and figure 8.11 we display the rate of multihoming under five different weighting schemes.⁴⁴ We first ask which weighting schemes make a difference, that is, which weights are correlated with the firm’s decision to multihome? For this purpose, the fifth column, “unweighted,” is the base case. On both base platforms, we see the same pattern. Column (2), weighting by “minutes per visitor,” predicts about the same rate of multi-

42. 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 (2007) 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.

43. Bresnahan and Reiss (1990).

44. The schemes are daily “visitors” (i.e., users), minutes per “visitor,” that is, 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 in the footnotes to Error! Reference source not found..

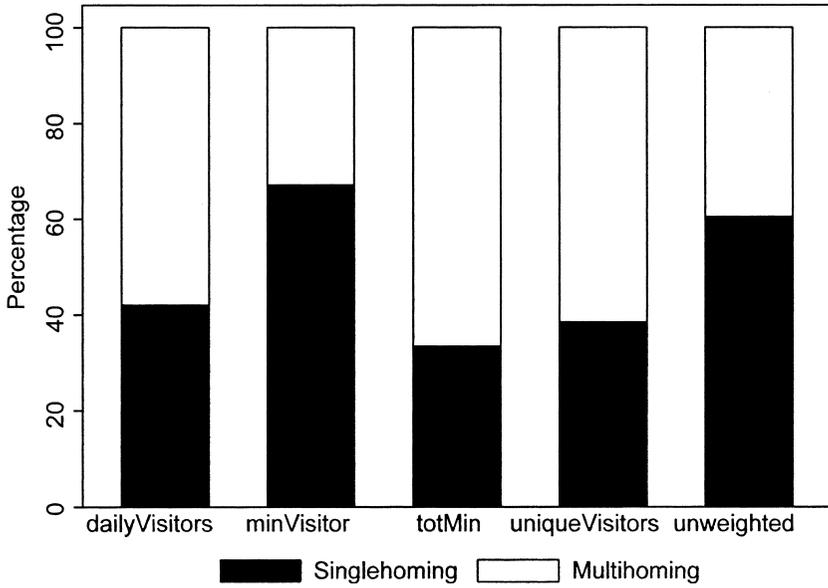


Fig. 8.10 Weighted multihoming rate on iOS

Source: comScore, January 2013.

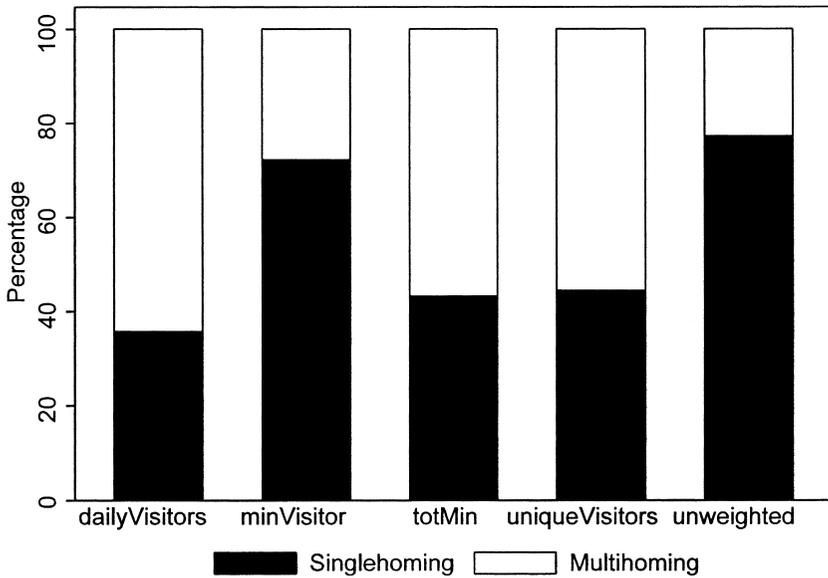


Fig. 8.11 Weighted multihoming rate on Android

Source: comScore, January 2013.

homing as the base case. The other three columns, unique monthly visitors, daily visitors, and total visitor minutes are all similar to one another and all, on both base platforms, predict a higher rate of multihoming than the base case.

The second fact is easy to interpret. When we weight by any of the three “size of the market” variables, we see more multihoming. That is, those developers who have a larger market size on the base platform are also likely to be observed on the other platform. This sounds very much like a familiar result from the entry literature. If (plausibly) market size on one platform is correlated with market size on the other platform, our weighting fact is no more than this: firms that expect a large market bear the technical and marketing costs of entering. If this is the right interpretation, even the weighted multihoming rates in columns (1), (3), and (4) are surprisingly low, another clue that the total fixed costs (technical plus marketing) of entering to serve consumers on a platform are substantial.

The other fact, based on the second bar (minutes per visitor), shows results much closer to the unweighted case. Now, industry figures are likely right that this measure is related to the intensity of preference at the individual user level—typically called “user engagement” in the industry. One might have thought that higher user engagement would mean more per-customer value for the app developer, and thus that there would be a significantly higher rate of multihoming than the base case. However, this does not appear to be the case. 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 certainly be the case for paid apps or subscription apps, or if high-engagement apps cannot show advertisements (profitably) twice as often if the app is open twice as long.

These weighting results establish clearly that more popular apps (larger market size) are more likely to multihome. The obvious interpretation is that a firm with a more popular product can pay the incremental fixed costs of supplying a second platform. Since the observable definition of multihoming in these results is appearance in both platforms on comScore, we cannot, by these results alone, distinguish marketing fixed costs from technical fixed costs. Other indicia, however, suggest that the marketing costs are high.

8.9.4 Multihoming at the Firm Level

One theory of platform choice frequently articulated in the early stages of the industry is the importance of firms’ technical capabilities. Some firms might have engineers trained in iOS or in the programming languages

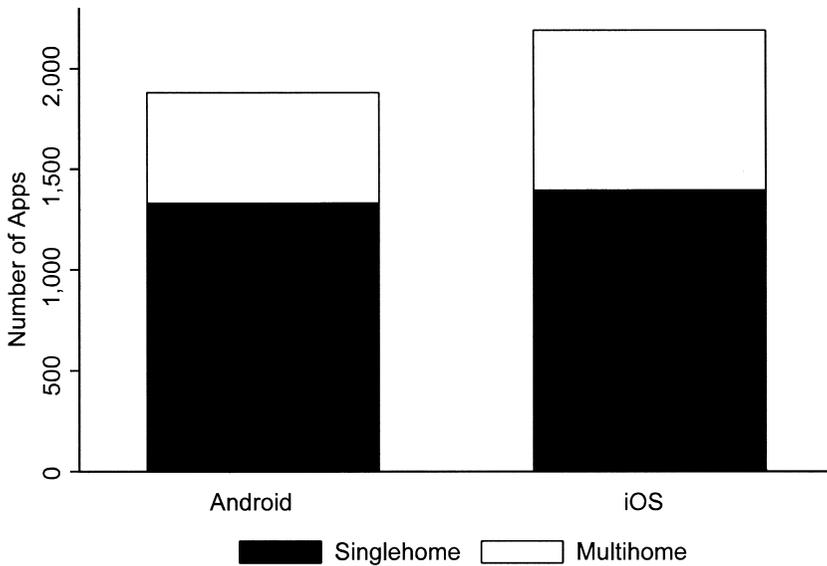


Fig. 8.12 Developer websites: This app available on the other platform

Source: App questionnaire, n: Android = 1,844, iOS = 2,117.

approved by Apple, while other firms might have the corresponding Android technical capabilities. We examine this hypothesis by looking at multihoming at the firm level.

We examine this hypothesis in figure 8.12 and figure 8.13. These figures are based on new samples, that is, the largest groups of apps for which we have observed developer websites. Once again, we use the “base” platform concept, that is, we examine the behavior of writing for the “other platform for all the apps we survey in the ‘base’ platform.” This sample goes considerably deeper into the long tail of apps on both platforms than does the sample of apps we were just examining. Also, the definition of multihoming here is writing an app for a second platform, not writing and marketing it to achieve success.

Figure 8.12 looks at multihoming at the app (product) level. The sample for the first column is the most popular apps on base platform Android, so “multihoming” means writing that same app for iOS. The second column is symmetric, an observation is a popular app on base platform iOS, and “multihoming” is writing that same app for Android. Figure 8.13 has the same structure and sample, but here “multihoming” refers to the same firm supplying any app to the other platform.

There are three interesting facts in these figures. First, the rate of multihoming at the app level (figure 8.12) has fallen in these broader samples compared to the data we saw in figure 8.9. Above, we had the most popular 1,231 apps and about a 64 percent rate of multihoming. Here, we are looking

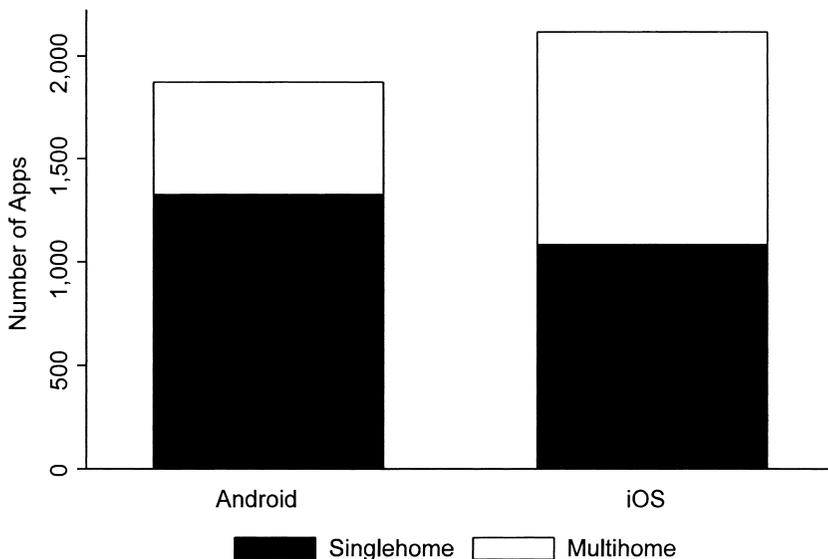


Fig. 8.13 Developer websites: Some apps available on the other platform

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

at the most popular 4,000 apps or so, and the rate of multihoming has fallen: 23 percent of Android apps and 29 percent of iOS apps are multihomed. Second, while the rate of multihoming has fallen, it has not fallen to zero. This encourages us, in future work, to examine the relationship between market size and multihoming more quantitatively to learn more about the incremental fixed costs of writing for a second platform.

The third fact comes from the contrast between figure 8.12 and figure 8.13. Simply put, there is much more multihoming at the firm level than there is at the app level. On both platforms, there is a considerable probability that a firm that does not write the same app for both platforms nonetheless has apps on both platforms. The difference is not small. Also, it is not symmetric. When the base platform is iOS, meaning that the firm has at least one popular app on iOS, there is a very large, 49 percent, probability that the firm will offer some app on Android.

This is strong evidence against a technical capabilities theory of platform choice. We see a large number of firms that already multihome today; they know how to develop for the other platform. They choose not to port all apps to the other platform; they are not constrained by lack of capabilities. This is interesting in itself about mobile firms; more importantly, it suggests market forces will be able to influence future developer platform choices should platform market structure move away from the current top-two-tiered structure.

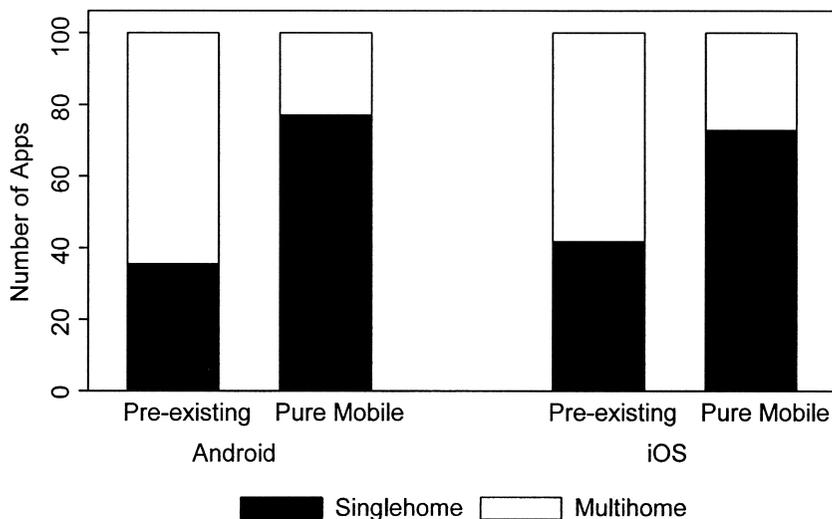


Fig. 8.14 Multihoming by firm scope

Source: App questionnaires. Based on the question: “Is the app produced by (or in partnership with) a larger company (one with a presence outside the mobile economy)?”

8.9.5 Multihoming by Firm Type

In the app questionnaire data set, we have also constructed variables for the type of firm associated with the app developer (purely mobile app firm or a preexisting, “corporate” firm offering nonmobile products). We report the rate of single homing and multihoming for each observable type of firm in figure 8.14. This question is based on our linking the developer websites to standard firm-information sources; linkage involves more hand work, as not all developers are all that specific about their parent firm.

The value of this linkage lies in investigating a hypothesis about the incidence of the marketing costs associated with the mobile industry’s market institutions for matching users to apps. A firm that was founded in the mobile industry, an entrepreneurial mobile app firm, must acquire customers on each platform using the mobile industry’s market institutions. A preexisting firm may already have customers, and can market its mobile app to them. Thus the marketing costs of porting to a second platform may be significantly lower for a corporate app than for an entrepreneurial app.

Figure 8.14 provides considerable evidence for the hypothesis that the costs of porting and multihoming *for entrepreneurial firms* 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, and so forth are, at 65 percent

(0.645), quite likely to multihome. For iOS, the probability of a top app multihoming if the developer is preexisting is 58 percent (0.58). These probabilities are two to three 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 our matching apps across data sets means that we do not have a fully integrated data set yet. Furthermore, we have not filtered out apps for whom the porting decision may not be interesting or may already be predetermined (wireless carriers, platform providers, etc.). A great deal of interesting analysis may arise from augmenting this data with information on categories (Davis, Muzyrya, and Yin 2013). Finally, we as yet do not have a serious model for understanding first and second platform choice by developers.

Still, after looking at a number of different indicia, we have reached what we think is an important interim conclusion. There appear to be substantial marketing costs of reaching both halves of the mass-market audience in the mobile industry, and these marketing costs fall far more heavily on entrepreneurial firms than on established consumer product and services firms. That suggests once again that market institutions may also be limiting the range of economic experiments about value creation. Another implication of our interim conclusion is that the market institutions of the young mobile industry are slowing the rate of innovation effectively delivered to consumers and biasing its direction toward corporate apps.

8.10 Alternative Equilibrium Scenarios

8.10.1 Competition among Platforms

The implications of all this developer behavior, including single homing, multihoming, and platform choice for the single homers for the competition between platforms is difficult to discern at this juncture. No tip to a single platform seems imminent, even though full 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.⁴⁵ There does not appear to be a strong expectation—in any direction—among consumers.

45. See, for example, Farrell and Klemperer (2006). 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.

To the extent that developers have expectations, those would appear to favor the original incumbent iOS/iTunes Store platform. (See table 8.1.) 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 and largely skipping the Blackberry or Windows Mobile efforts. As we saw in table 8.1, the count of larger developers in the com-Score data with Android-only exceeds the count with iOS-only, with a bit of a trend toward Android. For now, our interpretation is that the market is *not* 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. Media (books, newspapers, music, etc.) are significantly more likely to multihome than are app developers. One reason may be that media have buyer-seller matching market institutions that are outside the mobile economy and media have low technical costs of porting to a new platform. Media are natural multihomers.

Within media, however, new issues arise. 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. Indeed, some media 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 that 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, too, does the boundary between all these mobile devices—today primarily supporting consumption—and PCs—today primarily supporting work—in markets served. 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.

8.10.2 Divided Technical Leadership (DTL)

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.

Bresnahan and Greenstein (1999) argue that this structure has been historically important in platform competition in the computer industry; DTL was important in the early PC industry and in the fall of IBM’s hegemony in corporate computing. The structure has been important both in enabling new platforms to compete against an established platform and in enabling competition within a platform, even with positive feedback around a particular standard.⁴⁶ We have recently seen another very important example in the mobile space itself. Competition between mobile carriers such as AT&T and Verizon has been changed forever by the partially overlapping positive feedback loops around mobile devices and mobile platforms we are studying in this chapter. In light of its past importance, it is reasonable to examine DTL as one force that will affect future events in mobile platform competition.

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 that 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 is permitted to 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 offering “infrastructural” functionality, functionality that many

46. 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.

operating system providers would seek to offer only in the OS itself. Some hardware manufacturers, such as Samsung, have made impressive efforts to turn Android for mobile devices into their own platform, not Google's.

Apple set up the iOS/iTunes Store platform in a much more centralized and controlled way, and thus has the capability to 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 as easily. (Apple's enthusiasm for enforcing restrictions on developers may well have declined over time after the success of Android.) 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, 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.

While the current platform industry structure seems 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 and its extremely impressive capabilities as a "strong second" imitator, it seems certain that this firm will continue to invest in Windows Mobile for a long time (à 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 imagine 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 modified 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 (à la the shift from "the IBM PC and clones" to "the Windows PC").

8.11 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. In this regard, with the need for “mere marketing” to be successful for the industry to expand, mobile is like earlier rounds of ICT invention such as the mainframe computer, the superminicomputer, the PC, or the widely used Internet.

Second, the sheer size of the mobile opportunity illustrates an important change of twenty-first 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 great deal of experimentation can fit 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 new rounds of innovation by finding important general sources of value for mobile applications. This is an important difference from earlier ICT industries. 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 (<1 million) 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 industrial organization.

The 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 primarily the “localization” of applications first written in English; this time it means drawing on a global talent pool to create firms such as Rovio and Distimo.⁴⁷

47. 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 (1999) for a discussion of national versus international supply in the creation of a number of earlier computer markets.

Similarly, high device penetration has made it economically feasible in the mobile era to have app experiments that are very heterogeneous with regard to the developer's business model. What Hayek (1945) called the economic function of the entrepreneur, finding overlaps between supply opportunity and demand need, can be taken up by any organizational form—and it has been. Right now, the online app-discovery mechanisms tend to favor existing businesses that give their existing customers an app. That is, the entrepreneurial function is being taken up by firms that 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 remains 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 already 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 labeled “the problem of the tenses” three decades ago.⁴⁸

Why has the app exploration taken so long? This question falls into “the problem of the tenses”! A better question is: 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? This question leads us to 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, and so forth. 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

48. As we write, Google faces 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.

shop?⁴⁹ Even more difficult to foresee without experimentation is the market 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 (nonmobile) 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 percent 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 new “disruptive” 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, preexisting 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 made a 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, and Pandora are in competition with the music portion of the online stores.) Mobile tele-

49. 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 twenty-something app developers thinking about the concerns of twenty-something customers, in this case the incomprehensible tribal and mating behaviors of twenty-somethings. This is uncharitable; market experiments are heterogeneous because different experimenters have different knowledge, goals, and 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.

phone 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 app discovery 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 sixty 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 twenty-first 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.

References

- Armstrong, M., and J. Wright. 2007. "Two-Sided Markets, Competitive Bottlenecks and Exclusive Contracts." *Economic Theory* 32 (2): 353–80.
- Bar, F. 2001. "The Construction of Marketplace Architecture." In *Tracking a Transformation: E-Commerce and the Terms of Competition in Industries*, 27–49. Washington, DC: Brookings Institution Press.
- Bresnahan, T. F. 1992. "Sutton's 'Sunk Costs and Market Structure: Price Competition, Advertising, and the Evolution of Concentration.'" *RAND Journal of Economics* 23 (1): 137–52.
- Bresnahan, T. F., and S. Greenstein. 1999. "Technological Competition and the Structure of the Computer Industry." *Journal of Industrial Economics* 47 (1): 1–40.
- Bresnahan, T. F., and F. Malerba. 1999. "Industrial Dynamics and the Evolution of Firms' and Nations' Competitive Capabilities in the World Computer Industry." In *Sources of Industrial Leadership*, edited by D. Mowrey and R. Nelson. Cambridge: Cambridge University Press.
- Bresnahan, T. F., and P. C. Reiss. 1990. "Entry in Monopoly Markets." *Review of Economic Studies* 57 (4): 531–53.
- Bresnahan, T. F., and M. Trajtenberg. 1995. "General Purpose Technologies: 'Engines of Growth'?" *Journal of Econometrics* 65 (1): 83.

- Christensen, Clayton, Curtis W. Johnson, and Michael B. Horn. 2010. *Disrupting Class, Expanded Edition: How Disruptive Innovation Will Change the Way the World Learns*. New York: McGraw-Hill.
- Davis, J., Yulia Muzyrya, and Pai-Ling Yin. 2013. "Experimentation Strategies and Entrepreneurial Innovation: Inherited Market Differences from the iPhone Ecosystem." Working Paper, INSEAD.
- Farrell, J., and P. Klemperer. 2006. "Coordination and Lock-In: Competition with Switching Costs and Network Effects." Working Paper, Competition Policy Center, University of California, Berkeley.
- Greenstein, S. 2001. "Commercialization of the Internet: The Interaction of Public Policy and Private Choices or Why Introducing the Market Worked So Well." In *Innovation Policy and the Economy*, vol. 1, edited by Adam Jaffe, Josh Lerner, and Scott Stern, 151–86. Cambridge, MA: MIT Press.
- Hayek, F. A. 1945. "The Use of Knowledge in Society." *American Economic Review* 35 (4): 519–30.
- Helpman, E., and M. Trajtenberg. 1998. *Diffusion of General Purpose Technologies*. Cambridge, MA: MIT Press.
- Klepper, S. 2002. "Firm Survival and the Evolution of Oligopoly." *RAND Journal of Economics* 33 (1): 37–61.
- Rysman, M. 2009. "The Economics of Two-Sided Markets." *Journal of Economic Perspectives* 23 (3): 125–43.
- Sichel, D. E. 1997. *The Computer Revolution: An Economic Perspective*. Washington, DC: Brookings Institution Press.
- Sorenson, A. T. 2007. "Bestseller Lists and Product Variety." *Journal of Industrial Economics* 55 (4): 715–38.
- Sutton, J. 1991. *Sunk Costs and Market Structure: Price Competition, Advertising, and the Evolution of Concentration*. Cambridge, MA: MIT Press.
- Waldfogel, J. 2013. "And the Bands Played On: New Music in the Decade after Napster." Working Paper, Carlson School of Management and Department of Economics, University of Minnesota.