ABSTRACT

Title of Dissertation: MECHANISM DESIGNS TO MITIGATE
                          DISPARITIES IN ONLINE PLATFORMS:
                          EVIDENCE FROM EMPIRICAL STUDIES

                          Raveesh Kolake Mayya, Doctor of Philosophy,
                          2020

Dissertation directed by: **Professor Siva Viswanathan,**
                          Department of Decision, Operations and
                          Information Technologies

With the rising ubiquity of online platforms, there is an increasing focus on platforms' role in enabling fair exchanges between buyers and sellers. Traditionally, platforms have inbuilt mechanisms such as screening or upfront data-gathering disclosure that encourage transactions between unfamiliar participants. Since such mechanisms can introduce power disparities between different sides, platforms have enacted policy changes to fix the imbalance. Extant literature hasn't studied the unintended consequences of such policy changes. My dissertation seeks to fill this gap by examining platforms' decisions to enact policy/mechanism changes that level the playing field by decentralizing choices for different sides. Using empirical studies, my dissertation seeks to causally identify the impact of such changes on outcomes for participants as well as for the platform.

   The first essay in my dissertation examines the impact Airbnb's decision to make screening optional. There is increasing evidence that two-way screening mechanism has been used as a tool by users on the platform to discriminate against some users on the other side. In making screening optional, I find that African American hosts and female hosts are more likely to forgo screening

and they benefit the most (in terms of occupancy, price and/or ratings) from forgoing screening, indicating that making screening optional can serve as a useful mechanism in helping alleviate reverse discrimination of hosts by guests.

The second essay studies platforms' attempts to provide smartphone users with better choice over which sensitive information can mobile apps access. In particular, I examine the timing of mobile apps' decisions to upgrade to Android 6.0, which restricts the ability of mobile apps from seeking blanket permissions to sensitive user information at download, instead requiring them to request à la carte permissions at run-time. I find that apps that over-seek (access information that are non-essential to their functionality) sensitive information from users strategically delay upgrading to Android 6.0. However, these apps suffer popularity and reputational costs in the Android marketplace.

Collectively, the findings in my dissertation provides valuable theoretical as well as practical insights about the welfare implications of choice decentralization on all sides in online platforms, not just the intended side.

MECHANISM DESIGNS TO MITIGATE DISPARITIES IN ONLINE PLATFORMS:
EVIDENCE FROM EMPIRICAL STUDIES

by
RAVEESH KOLAKE MAYYA

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park, in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2020

Advisory Committee:
Professor Siva Viswanathan, Chair
Professor Peng Huang
Professor Rajshree Agarwal
Professor Guodong (Gordon) Gao
Professor Anna Alberini (Dean's Representative)

Acknowledgements

The five years at a Ph.D. program is the best chance for a budding scholar to learn how to identify high quality problems that are worth solving. During my Ph.D. journey, I have been fortunate to work with high quality researchers who have helped shape my research agenda and thinking process.

First and foremost, I would like to express my deepest gratitude to Professor Siva Viswanathan, the chair of my dissertation committee. He truly follows the philosophy of, "*give a man a fish and you feed him for a day; teach a man to fish and you feed him for a lifetime*" and puts in a phenomenal effort to ensure that we learn *fishing*, i.e., identify high quality research question. While he emphasizes that I identify my own research interests, he ensures that I do not stray or get fixated on inconsequential yet glamorous elements of a research exploration. He is a true philosopher in that I have learnt some of my best life lessons listening to his thoughts on mind, intelligence, action-reaction. I also sincerely want to thank Professor Peng Huang with whom I initiated my first research project at Maryland. I have thoroughly enjoyed my discussions with him about research, career, life after tenure, and the likes. Conversing with him over the years has helped me to be very careful and diligent about my research ideas and empirical methodologies. I also like to sincerely thank Professor Rajshree Agarwal. It is a pleasure to see her in action, both as a prolific researcher/thinker as well as a guide. There is a lot to learn from her about managing time and prioritizing tasks: she has managed her high research productivity while heading a very impactful, multi-disciplinary research center and guiding multiple PhD students, across the disciplines. My sincere thanks are also due to Professor Gordon

Table of Contents

List of Tables

Table for Chapter 2

Table for Chapter 3

List of Figures

Figure for Chapter 2

Figure for Chapter 3

Figure for Chapter 4

Chapter 1: Introduction and Overview

Over the past two decades, the platform economy has grown from being a blip in the radar of digital economies to its de-facto flag bearer. Today, the top 5 US-based technology platform economies have a market capitalization of over 5 trillion USD (Forbes 2020). With over 170 platform startups with greater than 1 billion USD valuation by 2016 (Accenture 2016; Wikipedia 2020), the platform economy is poised to be the most sought-after business model for startup ideas. One of the key reasons for such an explosive growth of platform economy is the digital platform's ability to provide efficient ways of discovering unfamiliar transacting partners and completing transactions (Eisenmann et al. 2006; Rochet and Tirole 2003), For example, Airbnb efficiently connects house owners/tenants that have spare inventory with others who are in need of a space to spend a night. In a way, the features of digital platforms democratized the access to markets that were held by institutional players until recently.

However, a platform's growth depends not just on bringing together different transactional partners and encouraging transactions between them but also on guaranteeing fair exchanges on both sides. Asymmetries between participating sides may result in consumers' platform abandonment, especially since competition among similar platforms have lowered the switching costs for consumers (Eisenmann et al. 2011; Patchin and Hinduja 2010). Therefore, it is in the platform's interest to detect and mitigate inadvertent disparities between its various transacting partners, thereby ensuring that both the sides are satisfied.

Platforms have traditionally focused on mitigating frictions and inefficiencies on platforms that hinder transactions between the different sides. As a result, platforms

have designed mechanisms such as screening or bundling privacy choices, which are aimed at encouraging and simplifying transactions between unfamiliar transacting partners in platforms. Most of such mechanisms have been designed to enable easy onboarding of initial set of users. For example, the ability to screen transacting partners allows uninformed participants in the transaction to ascertain quality before participating in the transaction (Stiglitz 1977). Especially when platforms such as Airbnb or Uber enables inventory sharing with unknown people, allowing users to screen seems a required mechanism. Similarly, allowing app users to easily scrutinize all the sensitive data that an app accesses before they decide whether or not they want to use the app is found to reduce the buyer's risk perceptions about sellers (Hui et al. 2007; Malhotra et al. 2004).

However, over time, platforms have realized that participants may not always utilize these mechanisms in a way that such mechanisms were originally intended to. Left to the market, some participants from one side of the platform may exploit these mechanisms to dominate the other side in transactions. Such a gaming of the system would result in imbalances on these platforms that would harm participants as well as the platform economy. For example, in online labor platforms, if platform design decisions end up shifting market power in favor of the job providers, the market may evolve to be inefficient and of low quality (Kingsley et al. 2015). Platform mechanism designed to mitigate asymmetric information or to reduce the cost of transaction may inadvertently favor the stronger, more established players among sellers, resulting in discrimination or strategic retaliation against the less advantaged participants on the platform (Edelman and Luca 2014; Klein et al. 2009; Ye et al. 2014). For example,

while blanket/bundled permission granting reduces consumers' cost of individually granting them, it ends up limiting their choices and makes it easy to a "valuable" app to bulldoze its way through over-accessing users' sensitive information. Some of the disparities have been so acute that they have attracted negative press and/or the attention of regulators (Levin 2017; Satariano 2019).

As a response to media coverages and regulator interventions, platforms have recently begun to identify and mitigate such potential disparities by introducing policy changes to improve user choices. Such policy changes are usually designed to shift this power imbalance towards the traditionally disadvantaged side on platforms. For example, when media/academia started highlighting the discrimination that guests faced on the platform, Airbnb in enacted multiple policies to prevent the guests being discriminated. However, given the multi-sided nature of the platforms, such policy changes will invariably have unintended consequences on all sides. Furthermore, given that many platforms are making policy changes, a systematic investigation of policy changes on all sides of the platform is imperative. However, my literature survey points to the lack of such systematic investigations into platforms' efforts to mitigate potential imbalances on platforms.

My dissertation seeks to address this gap in literature. Specifically, my dissertation essays focus on platforms' decisions to relax, alter or introduce novel mechanisms that are aimed at leveling the playing field for the different sides in these platforms. The two essays have focused on two distinct online platforms: an exchange platform and a smartphone platform. In each of the studies, I identify power-disparities faced by participants and the platform's effort to mitigate this problem by introducing new

mechanisms that increases users' choices. Interestingly, instead of forcing these mechanisms on all sellers, platforms introduced these market mechanisms as a choice to sellers on the platform. Sellers' choices here are with respect to whether and/or when they would adopt such policies (changes). Hence, my essays also investigate antecedents of decisions to adopt these mechanisms, in addition to examining their consequences. Then, I employ various well-accepted identification strategies to causally identify the impact of such mechanisms on outcomes for participants, as well as for the platform.

The first essay in my dissertation examines the impact of forgoing screening in an online sharing platform. Airbnb's rise in dominance has helped many house owners rent out their unutilized inventory for a price. Airbnb is a decentralized platform, in that the platform enables guests and hosts to obtain necessary information about each other and interact on the platform before mutually agreeing to transact. While this two-way screening is assumed an essential mechanism in a decentralized platform, it has also raised concerns of increased discriminatory or biased behaviors. Over a third of guest requests have been found to be rejected by hosts, resulting in market frictions (Fradkin 2015). Such excessive vetting also seems to systematically affect ethnic minorities (Cui et al. 2020), travelers with disabilities (Ameri et al. 2017), younger guests, or guests with children (Karlsson et al. 2017). In 2014, Airbnb introduced a feature called "Instant Book" which enables hosts to voluntarily forgo their ability to screen guests. Given that only 15 percent of the Airbnb hosts voluntarily opted to forgo screening of their guests, I first examine the antecedents of hosts' decision to forgo screening. Equipped with the knowledge of when a listing is more likely to forgo screening, I

causally identify the impact of adopting this mechanism on hosts' outcomes on the platform, specifically their occupancy, pricing and rating. I find that listing specific characteristics such as falling occupancy or listings with newer hosts are more likely to forgo screening. Interestingly, I find that African American hosts and women hosts are more likely to forgo screening. I explore and explain the potential reasons why such hosts may be more likely to forgo screening. Next, I find that the decision to forgo screening pays off economically, with a 13.52% increase in occupancy, an equivalent of an increase in revenue averaging $300 per month. However, such an increase comes at the cost of a marginal decline in review ratings (a 1.20% drop). I also find significant heterogeneity in race and gender regarding the economic benefits of forgoing screening. I find that African American hosts and female hosts benefit the most from forgoing screening. This study highlights that making screening optional may serve as a useful mechanism in helping alleviate reverse discrimination of hosts by guests. This study also highlights the importance of understanding the differential effects on mechanism designs based on race/ethnicity and gender.

The second essay studies the impact of Android's attempt to provide Android users with better control over their private information accessed by Android apps. In particular, I study apps' decisions to upgrade to Android 6.0, which restricts their ability to seek blanket permissions to sensitive user information at download, instead requiring them to request à la carte permissions at run-time. Prior to 6.0, the apps were required to upfront list all the permissions to access users' sensitive information before users even downloaded the app. While this mechanism reduced the cost of individually picking and choosing the permissions, the expectation was that app users would

download the app only if they agreed to grant all permissions. However, this hardly deterred apps from misusing this mechanism. Firstly, the information overload during the disclosure discouraged many users from even reading through the entire list before downloading the app. Even if the user reads the list and had to decide between (a) granting permissions to access patently non-essential information, and (b) not using the app at all, most users would choose the former, especially if the app is "valuable" in their mind or does not have an easy substitute. As a result of this choice structure, a significant proportion of apps in Android over-seek permissions (i.e., seek more permissions than those required for the app's functionality) (Felt et al. 2011) leading to a moral hazard problem. Such permission over-seeking behavior poses security and financial risks to users in the platform ecosystem (Sarma et al. 2012; Wei et al. 2012), resulting in the erosion of trust on the platform. With Android 6.0, Android allowed app users to (a) delay granting permissions until they start using the app and, (b) pick and choose which permissions to grant. What is interesting in this setup is that Android gave a three-year window for all the apps to upgrade (i.e., target the latest version of Android) instead of forcing all apps to upgrade immediately. Hence, in my setup, I observe apps' temporal choice of either upgrading to Android 6.0 immediately, which provides mobile apps with the latest platform features, or staying with an earlier version, which provides them with better access to users' sensitive information. By utilizing a unique panel dataset of 13,604 most popular apps for 24 months, I find that apps that traditionally over-seek permissions strategically delay upgrading. More specifically, I find that a key reason why such apps delay upgrading is to continue gathering such sensitive information to serve customized and targeted advertisement.

Interestingly, such decisions to strategically delay upgrade does not come without a consequence. I show that, while such strategic delays enable apps to continue accessing sensitive user information for a longer time period, such apps suffer popularity (a fall of 9.06% in app's popularity) and reputational costs (a net rating loss of about 560 users) in Android Marketplace by delaying upgrade. This research adds to literature on mechanism design by showing that carefully designing *when* (download time vs run-time) and *how* (blanket vs individual) users grant permissions improve sellers' information collection behavior and penalizes strategic sellers.

To summarize, a number of platforms are responding to hot issues (that attract media and regulators' attention) on the platform by introducing policy changes. When platforms, particularly the online markets, enact policies that enable individuals on the platform to use their own decentralized information, I find that there are welfare implications. At the same time, such welfare implications are not only for the side for which the policy changes are intended to. I find that the other side of the network is also strategic about such change in policies. Therefore, platforms definitely need to take the actions of other side into account while implementing policy changes. Collectively, the findings in my dissertation provides valuable theoretical as well as practical insights about the welfare implications of choice decentralization on all sides in online platforms, not just the intended side.

Chapter 2: Who Forgoes Screening in Online Markets and When? Evidence from Airbnb

2.1. <u>Introduction</u>

Screening—the vetting of potential trade partners—is a key mechanism to reduce information asymmetry frictions and alleviate the "lemons" problem (Akerlof 1970) in online peer-to-peer (P2P) markets (e.g., Horton 2017; Iyer et al. 2009; Lin et al. 2013). Screening helps determine the trade-worthiness of potential transactional partners by examining pertinent information (e.g. profiles, past transaction histories) or through direct communication, and to reject transaction requests from questionable parties. For example, passengers and drivers alike on Uber (a ridesharing platform) can view ratings before accepting a ride; as may chefs and guests on EatWith (a dining experience platform) or clients and workers on Upwork (an online labor market platform). Scholars note the increased market efficiency due to better matches (Barach 2015; Iyer et al. 2009), but also provide cautionary evidence of discriminatory/excessive screening by one party, say the provider, of the other party, say the client (Cui et al. 2020; Fradkin 2015; Kanoria and Saban 2017; Pallais 2014; Romanyuk 2016). For example, Cui et al. (2020) find that, in the absence of information, Airbnb hosts engage in discriminatory screening of guests resulting in guests with African American names being less likely to be accepted compared to requests from guests with White names. Chan and Wang (2017) show that clients in online labor platforms screen workers by gender and are biased towards hiring female workers.

Left unexamined, however, is the recent popularity of *voluntarily forgoing* the option to screen the other party, even though the other party continues to screen the focal party. For example, on the car-sharing platform Turo, car owners can enable the

"Instant Book" feature to directly confirm booking requests from travelers. Similarly, in online market Dogvacay for home dog boarding services, a dog sitter can enable "Instant Book" and automatically approve a guest's request without having to check the guest's profile and credibility. Despite its increasing usage, we lack systematic studies of the phenomenon. When providers forgo such screening, they are forsaking the benefits from verification of match quality and increasing the risk of encountering lemons. Why, when and who may choose to forgo screening, and what are the resultant outcomes of forgoing this option for themselves? I address this research gap in the setting of Airbnb—the online lodging marketplace platform—which expanded on its traditional mechanisms requiring hosts' approval of guests' booking requests (i.e., screening is enabled) by popularizing the "Instant Book" feature that allows hosts to automatically accept every booking request (i.e., screening is forgone) in 2014. The setting enables me to examine the voluntary decision of hosts to switch from traditional to instant booking, relate characteristics of market listings and hosts to such behavior, and investigate the market performance of the hosts' listings based on their behavior. Specifically, Figure 2.1 depicts my research framework and questions.

Given that most online P2P markets, and particularly my research context—Airbnb, initially offered only the traditional booking option, I treat traditional booking as the default norm (i.e., the *non-switchers)*. I ask and answer two research questions: the first question (RQ 1) aims to explore the antecedents of forgoing screening for a listing and asks, what characteristics of market listings (and the hosts) are associated with the likelihood of forgoing screening? The second question (RQ 2) aims for causal identification of the consequences of forgoing screening for a listing and asks, how do

9

the market outcomes—occupancy rate, listing price, and review ratings—for a listing change after it forgoes screening as compared to a listing that does not?

Based on a unique panel dataset of all Airbnb listings in New York City between August 2015 and February 2017, I conduct a variety of listing-level analyses to answer the proposed questions. I employ a logit model to explore what listing conditions and host characteristics are associated with a listing's likelihood of forgoing screening. My context represents a self-selected intervention, so I use the common approach of propensity score matching (PSM) in conjunction with difference-in-differences (DiD) analysis (e.g. Liu and Lynch 2011; Smith and Todd 2005) to causally identify the impacts of switching on a listing's market performance. Using this, I causally estimate the impact of forgoing screening on market outcomes through a simultaneous estimation system of three equations for a listing's occupancy rate, price, and ratings.

To preview my findings, forgoing screening is more likely when listings have lower demand, i.e., when listings have mid-level occupancies, face greater recent declines in occupancies, and face greater recent declines in price. This provides evidence that improving occupancy is one of the major drivers to voluntarily forgo screening. Switching to instant booking seems to pay off economically—I find occupancy levels increase by an average of 13.52 percent (or 1.79 nights per month) with no change in listing price, translating to an average of roughly $300 in increased revenue per month. However, this is at the cost of decreased ratings by about 1.20 % (or 1.08 points).

Underlying the above "average effects" based on demand conditions, I find significant heterogeneity in race and gender regarding who is more likely to switch, and their associated outcomes. African American hosts are more likely to forgo

screening than White hosts are, and they seem to reap higher benefits from doing so too. Forgoing screening results in African American hosts experiencing a higher percentage improvement in occupancy (18.37% as compared to 14.62% for White hosts), with lower declines in review ratings (a 0.28% decline as compared to 1.56% drop for White hosts). Turning to gender differences, my results show that female hosts are more likely to forgo screening than male hosts are, and they also gain more: listing price of a male host decreases by $4.28 relative to a female host; but there are no significant differences across gender in either occupancy rates or review ratings.

In addition to various robustness checks, I also perform falsification tests by analyzing Airbnb hosts' decision to switch back to screening. First, I find that switching back to screening is more likely for hosts whose listings experience greater drops in ratings from forgoing screening. Second, consistent with my result that African American hosts benefit more from increased occupancy from forgoing screening; I find that African American hosts are more likely to stay in the Instant Booking state relative to White hosts. While women are as likely to stay with Instant Booking as men are, African American female and male hosts are more likely to stay with Instant booking compared to white female and male hosts respectively.

Overall, my study highlights when, and for whom, forgoing screening may improve market outcomes. In being the first, to the best of my knowledge, to empirically examine the role of forgoing screening, my study makes important contributions to the literature on mechanism design in online markets and the literature on screening. While screening has been traditionally considered a necessity for online P2P markets to mitigate issues of information asymmetry, my study shows that forgoing screening can

benefit market participants in online markets. My study uncovers heterogeneity in response to, and benefits of having the option to choose or forgo screening; this enables online market participants to engage in their individual cost-benefit analysis and determine how to customize their choices based on demand conditions, and their own characteristics. My study also contributes to the emerging literature on the sharing economy by shedding light on mechanism designs in online markets and online sharing platforms. While several studies have focused on discrimination of African American guests by hosts in online platforms, my study shows that flexibility and options for screening can particularly benefit African American and female hosts, who traditionally have lower average occupancies and/or price compared to other hosts. In doing so, I highlight the importance of understanding the differential effects on race/ethnicity and gender in examining the impacts of mechanism design in online markets.

2.2 <u>Theoretical Background</u>

As backdrop and motivation for my study, I first provide below a focused literature review related to two key features of online P2P markets: screening mechanisms for addressing information asymmetry, and within-heterogeneity in demand conditions facing individual listings (products or services). I then discuss the antecedents and consequences of forgoing screening.

2.2.1 Screening as a Mechanism for Addressing Information Asymmetry in Online Markets

P2P markets are characterized by decentralized transactions, with little opportunity for direct contact or extended communication; moreover, most transactions are one-time interactions among market participants. Accordingly, information asymmetry is a

particularly thorny issue as both sellers and buyers encounter problems of adverse selection and moral hazard (Akerlof 1970), and screening is one of the key mechanisms to alleviate it.[1]

Screening permits an uninformed participant agent to take actions to ascertain the quality or type of the potential trade partner involved in a transaction (Stiglitz 1977). Two mechanisms of screening under imperfect information are identified in previous studies: *indirect screening* and *direct screening* (Buchmueller 1995; Hoff and Stiglitz 1990). *Indirect screening* refers to the uninformed agent's practice of offering a menu of contracts or options to the trade partner to make her self-reveal her type (Bastani et al. 2015). In *direct screening*, the uninformed agent exerts efforts to actively gather, investigate and evaluate information about the trade partner to infer her type (Cornell and Welch 1996; Bar-Issac and Cunat 2014). In this study, I focus on *direct screening* which is more often observed in online platforms. Within online P2P markets, direct screening mechanisms to help infer the hidden quality of a transactional partner include the ability to review and verify potential partner's information (e.g. Facebook profiles), track and analyze past transactional histories, or direct communicate via messaging and chat. For example, Airbnb provides various tools for hosts to know a guest, including accessing reviews on the guest from other Airbnb members, checking if the guest has a verified ID or other identity such as Facebook profile, and private messaging that allows hosts to infer the quality of the guest based on textual communication.[2] From the host's perspective, these direct screening tools help them avoid lemon guests that

---

[1] Signaling, which permits high quality participants to take actions to signal their quality and thus separate themselves from the lemons, is another mechanism to mitigate information asymmetry (Dewally and Ederington 2006).

[2] https://www.airbnb.com/help/article/1308/does-airbnb-perform-background-checks-on-members

could cause various forms of damages, such as leaving a mess that requires extra cleaning from hosts, ignoring the housing rule and causing discomfort to the hosts, and leaving unexpected bad views that could hurt future occupancies. Direct screening also helps the hosts avoid inconvenience and intangible costs that some guests may cause.[3] For the sake of brevity, hereinafter I refer to *direct screening* as screening.

Researchers have documented the use and benefits of screening mechanisms in online labor markets where client screening results in higher quality workers and output at lower prices (Barach 2015), and in online lending markets where lenders' screening of borrowers based on peer group comparisons mitigates adverse selection (Weiss et al. 2010), or lender discernment enables the provision of lower interest rates to borrowers who are better credit risks (Iyer et al. 2009). However, scholars have cautioned about screening also leading to excessive vetting, resulting in inefficient matching between the seller and the buyer. Fradkin (2015) documents increase in market frictions on Airbnb because almost a third of buyer requests are potentially rejected due to the hosts' screening. Pallais (2014) shows that inefficient matches between clients and workers in online labor market Upwork (used to be called oDesk) can be improved if the clients disregard screening and offer inexperienced workers a first job. Analytically, both Romanyuk (2016) and Kanoria and Saban (2017) show welfare gains from reduction of excessive screening and provide scenarios where some ignorance about the buyer's quality from the seller side may improve the performance of matching markets. On a related note, scholars have also documented, from the

---

[3] Even though Airbnb provides insurances for hosts against potential vandalism of their properties, it is the host's responsibility to prove that the damage is caused by the guest. In many cases, it is a time-consuming process. Such insurances do not cover for nuisance and annoyances that guests may cause to hosts and/or other guests sharing the premises.

perspective of the party being screened, the negative effects of screening in the form of discrimination. For examples, researchers have documented discrimination against travelers with disability (Ameri et al. 2017), younger guests and guests with children (Karlsson et al. 2017).

Thus, existing research depicts mixed findings for the use of screening mechanisms; while scholars note the beneficial effects of screening for the agent engaging in this mechanism, there is evidence against its welfare-enhancing role, particularly when taking into account the increased discrimination against the party being screened.

2.2.2 Heterogeneity in Demand Conditions for Listings in P2P Markets

Even under conditions of perfect information, there is significant heterogeneity in price and quality combinations in transactions in P2P markets. Accordingly, in addition to overall demand conditions facing *all* providers of an online listing (product or service), individual listings may differ in whether they are in higher or lower demand, and such heterogeneity relates to real or perceived differences in quality and price relative to alternative options (Dewan and Hsu 2004; Gunter and Onder 2017; Jin and Kato 2006). For example, within lodging sharing platforms such as Airbnb—my research context— while "average" demand for rentals is location and time specific, the characteristics of the individual listings—the listing price, amenities, geographical location, etc.—imply variation in the specific demand conditions faced by the focal provider.

In addition to the characteristics of the listing, individual demand could also be influenced by host characteristics such as race/ethnicity and gender. Such discrimination (statistical or taste based) has been well documented in various studies. Studies of labor markets have also found a preference for hiring Whites over hiring

African Americans (e.g., Pager et al. 2009; Reimers 1983) resulting in persistent racial wage gaps (e.g., Black et al. 2006; Smith 1993). Other studies document African Americans receive less favorable treatment—fewer rewards for the same level of "good" but more punishments for the same level of "bad"— than Whites (e.g., DeSante 2013; Steffensmeier and Demuth 2006). An underlying driving factor of such discrimination in markets with information asymmetry relates to racialized bias that Whites are perceived as more trustworthy and honest than their African American counterparts (Doleac and Stein 2013; Mays et al. 2007). In P2P markets specifically, studies of baseball card trades note offers received by African American agents were inferior relative to Whites agents at shows (List 2004); with 20% reduced prices observed in a field experiment on eBay (Ayres et al. 2015).

Studies of online lodging platforms reveal similar lower prices for equivalent listings by African American hosts relative to White hosts (Edelman et al. 2017; Kakar et al. 2017; Wang et al. 2015). Within P2P lending platforms, African American borrowers are less likely to receive funding (Pope and Sydnor 2011). Discrimination can also occur in the other direction (sellers discriminating against buyers). For example, Cui et al. (2020) conduct two randomized field experiments on Airbnb; their findings reveal racial discrimination by hosts who infer guest quality based on race/ethnicity, resulting in lower numbers of requests accepted from African American guest names relative to requests from guests with White names.

Similarly, gender-based prejudice has also been reported in a variety of contexts. Gender gaps in economic outcomes, including wage differentials, are consistently reported despite the increasing participation of women in the labor force and the

16

strengthening of women's education background and economic decision-making abilities (e.g., Altonji and Blank 1999; Bobbitt-Zeher 2007). Prior field studies have also found evidence for price discrimination against women as they may be perceived to have lower reservation prices (Ayres and Siegelman 1995).

In P2P markets specifically, Kricheli-Katz and Regev (2016) analyze over a million transactions on eBay to find equally qualified women sellers received fewer bids and about 80 cents for every dollar when selling identical new products relative to male sellers. Within P2P labor market platforms, male workers are less likely to be shortlisted and hired by employers (Chan and Wang 2017). Similarly, Stroube (2017)'s study of online lending platforms utilizes an exogenous policy shock to reveal taste-based bias against women borrowers.

Thus, in addition to information asymmetry, another important characteristic of online P2P markets is the heterogeneity of demand conditions for listings, whether driven by listing characteristics or provider characteristics.

2.2.3 Antecedents and Consequences of Forgoing Screening

Several gaps emerge from the above literature review. One, the literature on screening mechanisms leaves uncontested the implicit assumption that screening is always beneficial from the perspective of the party who engages in screening. This runs counter to the growing popularity and voluntary willingness to *forgo* screening in online markets. Many online markets, as noted in the introduction, enable an "Instant Book" feature, where sellers can select *not* to engage in screening of buyers. Two, the literature on solutions to information asymmetry is not fully integrated with the literature on heterogeneity in providers and buyers, even in the presence of full information. These studies do not account for the impact of within-heterogeneity of

17

listings on differential use of screening mechanisms among providers. Accordingly, studies examining screening mechanisms in P2P markets focus on costs and benefits of these on overall market efficiency, or for the *other* party subject to screening from the focal party. To the best of my knowledge, there are no studies that examine whether forgoing screening can benefit the party choosing to forgo screening, or when and to whom such benefits may be greater.

In this study, I seek to answer *when* and *for whom* I am likely to observe forgoing of screening, and *what* are the outcomes of doing so. At the outset, I note several boundary conditions. First, I am interested in markets where providers have the option of both retaining screening (traditional booking), and forgoing it (instant booking), and there are low costs of switching, so providers can switch between maintaining and forgoing screening relatively easily. Second, I focus on the focal provider's perspective, rather than the perspective of market efficiency, or the other party's perspective. Third, I focus on online markets where the other party can continue to screen, i.e., the focal provider is still subject to potential buyers screening and ascertaining their own trade-worthiness.

My first research question examines the antecedents of forgoing screening in an exploratory manner, seeking to uncover the factors associated with the likelihood of forgoing screening. From a focal provider's perspective, forgoing screening offers the benefits of a potential increase in the demand for their listings, because the "Instant Book" feature accommodates buyer time and convenience preferences, and may appeal to buyers who are in the new market participant segment. However, forgoing screening implies an increased risk, and thus cost, of having more potential lower quality buyers.

Additionally, lower quality matches absent screening may result in lower ratings or reviews, which has the cost of potential lower demand in the future. Thus, when deciding on whether to forgo screening, sellers have to assess the benefits and costs relative to retaining the option, as it relates to their own specific listing. It stands to reason then, that the use of screening mechanisms will have different associated costs and benefits for different listings. For listings in demand, the benefits of market expansion from forgoing screening are limited. However, for listings with over-capacity or excess holdings, the benefits of increased demand from a larger pool of guests as a consequence of forgoing screening could outweigh the potential increase in the costs. Based on this cost-benefit calculus logic, I expect listings that experience lower (or declining) individual demand to have higher excess capacity, and thus are more likely to forgo screening.

Equipped with the knowledge of *when* a listing is more likely to forgo screening, my second research question examines the market outcomes associated with forgoing screening in a causal manner. Again, consistent with the cost-benefit logic, I expect outcomes to be consistent with a "separating equilibrium." The market outcomes for listings and providers will reflect the commensurate benefits and costs, given endogenous decision making by providers. I expect that immediate returns—in terms of higher occupancy rate would be higher for switchers relative to comparable non-switchers, but this would be accompanied by lower reviewer ratings and/or lower listing prices, reflective of higher costs too.

As the characteristics of the provider in addition to the characteristics of the listing also contribute to the differences in demand conditions they face, the cost-benefit

calculus will differ based on these factors. For example, considering that African American hosts and female providers face discrimination and consequently lower demand in online markets as suggested in prior literature (e.g., Edelman et al. 2017; Kakar et al. 2017), they may have more incentives to forgo screening to benefit from a larger pool of guests. Therefore, in answering the two research questions, I also examine how host gender and race/ethnicity may play a role.

2.3 **Research Context and Data**

I examine the above research questions for the Airbnb platform, the world's largest online marketplace for short-term rentals (WSJ 2017). The platform allows house owners (i.e. hosts) with unused space to rent them out to guests who seek accommodations. Launched in 2008 in USA, Airbnb now operates in more than 65,000 cities in over 190 countries, with greater than 3 million active listings.[4] While initially limited to traditional booking, Airbnb introduced the "Instant Book" option in 2014. The traditional booking option represents two-way screening: the first step is where a guest screen across listings and submit a request to the host, the second step is where the host screens and either accepts or denies the request. With "instant booking", hosts voluntarily forgoing their own screening of the guest and eliminate the second step in the above process.

Several elements of this research context make it ideal for my study. First, Airbnb is a decentralized P2P platform with participants (hosts and guests) making their own choices. This allows me to study the impact of market mechanisms such as "Instant Booking" on individual outcomes. Further, switching to "Instant Booking" is relatively

---

[4] See https://www.airbnb.com/about/about-us

simple and seamless on Airbnb as hosts can set house rules, guest requirements, and length of stay, when choosing the "Instant Book" feature. Third, it is worth noting that the "Instant Book" feature requires only the hosts to waive their rights to screen their guests but allows the guests to continue to search through the listings, screen them and shortlist the ones they prefer. Finally, the context suits my data needs: as described in greater detail below, the analysis requires data about listing and host characteristics, which not only determine their screening choice, but also are also visible to potential guests for screening. Additionally, it requires data on outcomes associated with choosing to implement or forgo screening, such as subsequent occupancy, price, and ratings.

I compile the data by triangulating across publicly available data on Airbnb and proprietary data procured from a business intelligence firm that publishes monthly reports on Airbnb for over 100 cities. My sample consists of all Airbnb listings in the New York City (NYC) between August 2015 and February 2017. To ensure generalizable implications for a typical and active listing on Airbnb, I exclude listings priced above $1,000 a night at any month during the observation window, or listings who appear for less than a year[5], resulting in 13,757 listings in my panel. It is pertinent to note that most listings (about 96%) switch to instant booking at most once; only less than 4% of all listings switch to instant booking twice or more. This suggests that forgoing screening is not a period-by-period decision for most hosts. Instead, hosts prefer to stay with either traditional booking or instant booking, whichever works better for them based on their individual cost-benefit calculus. Consistent with my research

---

[5] Listings that are priced at over $1000 per night are very rare and often accommodate exceptional guests with special needs (e.g., party). Therefore, we exclude such atypical listings from my sample.

questions, the data permit me to differentiate across two groups of listings: (1) listings that always stay in the traditional booking status; (2) listings that are initially in traditional booking, but later switch to instant booking (i.e., switchers). These two categories constitute 96% of all listings; traditional booking represents the dominant choice, and instant booking represents 16% of the listings in my sample. While my definition of switchers includes listings that switch back to traditional booking later on, I only include their observations prior to switching back in the main analyses while ensuring robustness of my results in additional tests without such listings. Particularly of note is the very small count of 533 listings that forgo screening right from the onset. While I exclude these observations in the main analyses, I ensure robustness of my findings by including them in additional tests. My final data as described above, results in 196,155 listing-month observations in the main analyses. For each listing, I have time varying information about its characteristics (e.g., cancellation policy, room type, number of photos, etc.), as well as characteristics of the host (gender/race, Superhost designation[6], year of entry into Airbnb, displayed name, displayed picture, etc.). I also have information on the local competition (e.g., number of competing listings in the same neighborhood, number of Superhosts in the same neighborhood, etc.), and data on the listing's monthly outcomes (listing price, occupancy rate, review ratings).

2.3.1   Variable Definitions

Table 2.1 provides a detailed description of the variables used in the major analysis. The summary statistics and the correlation matrix are provided in Appendix Tables A1 and A2.

---

[6] On Airbnb, hosts who honor their booking reservations with at least 10 accommodations, 90% response rate, and 80% 5-star reviews are list as Superhosts (see https://www.airbnb.com/superhost)

Dependent Variables: The first dependent variable—the choice of forgoing screening is coded as a dummy variable, which takes the value of 1 if the listing is available for instant booking, and 0 if the listing has a traditional booking status. The other three variables measuring outcomes relate to the occupancy rate, price and the reviewer rating of the listing. I measure the occupancy ratio as the fraction of days the listing was booked (occupied) relative to the number of days the listing was made available in Airbnb. The price is based on the listing price in US dollars for the listing-month observation, and reviewer ratings are obtained for each month for each of the listings in Airbnb. On average, a listing is made available for booking 24.5 days per month, its listing price is \$167.87, and its monthly occupancy rate is 0.54 (or 0.54*24.5=13.23 nights).

Independent Variables: The key independent variables include listing specific variables representing its demand conditions, and host specific variables. In particular, prior levels and changes in occupancy rates, price and rating reflect the demand conditions specific to the listing, after accounting for many controls as listed below and in Table 1. The host specific variables include dummies that indicate if the host is a Superhost, or if the host is professional. [8] I measure host tenure as the number of days since the host has registered on Airbnb. For each host, I determine the race/ethnicity and gender as described in the section below. I include variables that capture local competition by using the similar listings panel present for each of the listings (12 similar listings per panel). Finally, listing specific control variables include review count, cancellation policy, type of listing (one of shared room, private room or entire property), minimum nights that a guest needs to book the place for, count of photos,

and popularity of listing (measured as the count of competition listings in which the focal listing appears in the similar listings panel).

Chief among the host variables are the race/ethnicity and gender. I ensured accuracy of race/ethnicity, and gender classification of the hosts based on inspection of profiles displayed on Airbnb. The race/ethnicity categories include White, African American, Asian, Hispanic, and uncertain (e.g., the host name is uninformative such as A.J and the profile picture is obscure and only shows the back of a person). The three gender categories include male, female, and uncertain (e.g., the host name is in abbreviations such as D.J. and the profile picture does not show any face). I coded the data manually, given the advantage of the manual coding approach over the algorithmic approach for identifying stylized names (e.g., @m@nd@ instead of Amanda) and low-quality photos (e.g., caricatured/tilted/side profile photos). Further, I validated the accuracy of the manually coded host race/ethnicity and gender labels by randomly selecting 20% of the sample and cross-validating the labels with those predicted by computer algorithms (Ambekar et al. 2009; Chan and Wang 2017; Yang et al. 2006). I combined two algorithmic approaches when doing so, the first relies on predictive models based on host names[7], and the second relies on predictions using host pictures through facial recognition algorithms.[8] Neither algorithm is necessarily superior to the

---

[7] Algorithmic approaches to predict host race/ethnicity and gender based on host names infer probabilities from classifiers based on Hidden Markov Models and decision trees (Ambekar et al. 2009). For race/ethnicity, consistent with Pool et al. (2015), I assign a race/ethnicity category to a host if the probability of the host name belonging to that category is at least 85% and mark the host race/ethnicity as uncertain if otherwise. I find no occurrence of cases of categorizing a name as belonging to two or more races/ethnicities (i.e., the probability of a name belonging to two races/ethnicities is at least 85%) in my dataset. For gender, following prior studies (e.g., Chan and Wang 2017; Tang et al. 2011), I match the host's first name to a database of first names with annotated genders constructed from public Facebook profile pages of 1.67 million users in NYC to algorithmically identify the gender of a host.

[8] Algorithmic approaches to infer host race/ethnicity and gender based on extracted facial features (e.g., eyebrow thickness, hair color, etc.) use profile pictures and employ computer vision and machine

other for revealing information about race/ethnicity and gender, which is why I combined both and designate the race/ethnicity and gender of a host as algorithmically identified if the labels generated by the two algorithms match and are marked as uncertain if otherwise. My cross-validation between manually coded labels and algorithm generated labels yields an agreement of 93.75% for host race/ethnicity and an agreement of 99.46% for host gender. These high levels of agreements provide confidence in my use of manual coding; in additional robustness tests, I ensure the findings are consistent across each type of approach (manual coding; name-based inference and picture-based inference) and their combinations.

Table 2.2 shows the distributions of host race/ethnicity and host gender in my sample. 55.73% of the hosts are White, and 12.44% of the hosts are identified as African American. This distribution is consistent with prior studies which show that Airbnb is predominated by White hosts overall (e.g., Edelman et al. 2017), and in New York City neighborhoods (Cox 2017). The ratio of female hosts to male hosts based on identified gender is about 53:47. Again, this distribution is fairly consistent with prior studies documenting that females account for roughly 55% of the Airbnb host community (e.g., Airbnb 2017; Mohlmann 2015).

2.4 <u>Analyses and Results</u>

2.4.1 When and Who Chooses to Forgo Screening?

To examine the factors that influence the voluntary choice of forgoing screening—my first research question—I use the following logit model:

---

learning techniques. Following prior studies (e.g. Chan and Wang 2017; Yang et al. 2006), I utilize a neural network trained facial recognition API to infer the race/ethnicity and gender of a host.

$$logit(\textbf{\textit{instant\_bookable}}_{it}) = c_0 + c_1 * occupancy_{it-1} + c_2 * occupancy^2_{it-1} +$$
$$c_3 * \Delta occupancy_{it-1} + c_4 * price_{it-1} + c_5 * rating_{it-1} + c_6 * host\_controls_{it} + \quad (2.1)$$
$$c_7 * listing\_controls_{it} + c_8 * competitor\_controls_{it} + c_9 * month\_dummy_t + \varepsilon_{it}$$

where the subscript i indexes the listings and the subscript t indexes the months.

Behaviorally, equation (2.1) models a listing's decision of switching from the initial traditional booking status to instant booking status during my observation period. As noted earlier, despite that in theory a host could forgo screening at any given time, in practice most listings do not frequently switch back and forth between screening and forgoing screening. Instead, they switch to instant booking at most once. In this case, the choice of forgoing screening is irrelevant once a listing has already switched. Therefore, following prior research (e.g., Gopal and Gao 2009; King et al. 2005), I drop observations related to a listing from the analysis after the listing has switched to instant booking. As argued earlier, the cost-benefit calculus of forgoing screening for a listing is related to the listing's demand situation such that lower (or declining) demand is associated with a higher likelihood of forgoing screening due to a higher chance that the received benefits of increased occupancy can outweigh the cost of reduced quality of guests. Therefore, I model a listing's decision to switch to instant booking in month t to be affected by its occupancy in the previous month t-1 and its recent change in occupancy from month t-2 to month t-1. I also include a quadratic form of $occupancy_{it-1}$ to detect possible curvilinear relationship between $occupancy_{it-1}$ and the propensity to switch to instant booking. In addition, I control for the other market performance indicators of the listing such as price and review rating in the previous month and their recent changes. I note that while my main analysis focuses on the prior period, additional robustness checks confirm results are qualitatively similar when the levels

and changes are computed in the prior periods ranging from three to six months. My choice of the prior period is to minimize the loss of observations due to differencing. I control for several variables that might also drive a host's decision to forgo screening, such as: (1) host tenure—novice hosts may use instant booking to learn to evaluate the quality of guests, (2) if the host is a professional host—professional hosts like hotels are used to the practice of not screening guests, (3) number of competitors in the neighborhood that have forgone screening—competitive pressure may force a host to mimic her competitors, and (4) minimum number of stay—a listing with a higher threshold for minimum number of stay may be more concerned about possible lemon guests. I also control for seasonality by including month fixed effects in the model. Given my interest in exploring how host gender and race/ethnicity may be associated with the decision to forgo screening, I do not include host fixed effects.

The results are shown in Table 2.3. The coefficients of $occupancy_{it-1}$ is significant and positive, and the coefficient of the quadratic form of $occupancy_{it-1}$ is significant and negative. These results suggest that having mid-level occupancy levels is positively associated with the likelihood of forgoing screening for a listing. The coefficient of $\Delta occupancy_{it-1}$ is significant and negative, indicating that experiencing recent declines in occupancies is positively associated with the likelihood of forgoing screening for a listing. Similarly, the coefficient of $\Delta price_{it-1}$ is significant and negative, suggesting that there is a positive association between experiencing recent price declines and the likelihood of forgoing screening for a listing.

In terms of host-related controls, I find that being a "professional host" (i.e., a host with three or more listings in NYC, typically a traditional bed and breakfast place)

is positively associated with the likelihood of forgoing screening. I also find being relatively new to the platform or non-Superhosts has a positive association with the likelihood of forgoing screening, as seen by the negative and significant coefficients of $host\_tenure_{it}$ and $host\_is\_superhost_{it}$. I also find that having a higher threshold for minimum number of stays and thereby greater concern about potential lemon guests is positively associated with the likelihood of forgoing screening, as indicated by the negative and significant coefficient of min_stay_{it}. [9] In terms of competitor-related controls, I obtain a positive & significant coefficient of $competitor\_instant\_bookable_{it}$.

I present the full model containing the effects of host race/ethnicity and host gender in column (2) of Table 2.3. I find African American hosts have a higher likelihood of forgoing screening as indicated by the significant and positive coefficient of $African$ $American_i$, while White hosts are less likely to forgo screening as indicated by the significant and negative coefficient of $White_i$. In addition, female hosts have a higher likelihood of forgoing screening than male hosts as indicated by the significant and negative coefficient of $male_i$. I also conduct additional sub-sample analyses (see Appendix Table A3) and find that among females, African American females are more likely to forgo screening than White females and among males, African American males are more likely to forgo screening than White males. Finally, in unreported tests, I conduct two additional tests using alternative estimation models and samples: a linear probability model, and by including observations of listings after forgoing screening and treating the decision as a period-to-period decision. I obtain consistent results.

---

[9] In additional analyses, I find that: (1) overall a host does not change *minimum_stay* in 86.76% of the months; (2) listings with lower occupancy rates are not different from listings with higher occupancy rates in terms of adjusting *minimum_stay*. These results help rule out the alternative explanation that hosts adjust *minimum_stay* as a strategic response to lower occupancy rates.

2.4.2   Impacts of Forgoing Screening

2.4.2.1 Difference-in-Differences Estimation Model

To answer my second research question examining the impact of forgoing screening on a listing's market performance, I adopt a difference-in-differences (DiD) model, which is commonly used to infer the causal effect of a "treatment" on the treated (Meyer 1995). In my case, the *treated* are the listings that have switched to instant booking (as compared to the control group of listings that stay in the traditional booking mode). Given the treatment in my context is staggered as each treated listing may switch at different times, I follow prior research (e.g., Autor 2003; Fang et al. 2014; Gao and Zhang 2016) and normalize the time dimension: For each treated listing, I set the month it received treatment to time 0, with months in the pre-treatment period in reverse chronological order denoted as month -1, -2, etc. and months in the post-treatment period sequentially denoted as month 1, 2, etc.

There are three key threats to inference to address in my DiD estimation approach. First, who receives the treatment is non-random, as discussed above in my first research question, so I need to address the selection bias in who receives the treatment. Approaches to addressing this issue include matching a control group with the treated group (Abadie 2005), using matching methods such as propensity score matching, which have been employed in various non-experimental settings when the assignment of treatment is not controlled by the researcher (Dehejia and Wahba 2002). Propensity score matching (PSM) has been frequently used in combination with DiD analysis for causal inference of the effect of a non-exogenous intervention (e.g. Liu and Lynch 2011; Smith and Todd 2005). Therefore, following the PSM method and taking into account the dynamic timing of treatment across listings, I first predict the

propensity score of a listing receiving the treatment at the observed treatment month (e.g., June 2016) using a logit regression on key observable covariates (i.e., time-variant covariates such as occupancy level, listing price, review rating, host tenure and number of reviews, and time-invariant covariates such as listing type and listing location) in the month prior to the treatment (e.g., May 2016).[10] I also ensure each matched listing is within a three-mile radius of the treated listing (calculated based on the zip codes of listings) to further eliminate any difference caused by geographic areas. To minimize the bias in the estimated treatment effect, for every treated listing I apply one-to-one nearest neighbor matching without replacement on the propensity score to identify a matched control listing that is used as a match only once (Austin 2010).[11] A covariate balance check shown in Appendix Table A1 provides evidence that a treated listing and a control listing are similar in all covariates after PSM despite their significant differences before matching.

I further check if the distributions of host race/ethnicity and of host gender remain consistent after the propensity score matching. As shown in Table 2.4, the pre-matching distribution of host race/ethnicity is comparable to the post-matching distribution of host race/ethnicity. Similarly, the distribution of host gender remains qualitatively unchanged after the matching. These results suggest that the sample after matching is

---

[10] I additionally perform two additional robustness checks to ensure that my results are not biased by the choice of covariates in PSM. First, I use the historical average of non-cumulative time-variant covariates (i.e., listing price, listing occupancy, number/count of ratings per month) in all months prior to the treatment month for the propensity score matching. Second, we introduce more covariates in the propensity score matching, including host race/ethnicity, host gender, number of listing photos, listing popularity, the host's Superhost status, and if the host is a professional host. In both cases, even though the matched sample of treated listings and control listings are different, we obtain consistent results in all the major analyses.

[11] The common support condition is validated and enforced as a significant overlap of propensity scores between the treated group and the untreated group is observed.

representative of hosts of various races/ethnicities and genders in the overall sample. Furthermore, treated and control listings are similar in terms of host race/ethnicity distribution and host gender distribution, providing further confidence on the validity of the sample used in the analysis. The matching procedure results in a sample of 1945 treated listing out of 2314 available treated listings and 1945 matched listings in the control group that is used for the DiD model as shown below. [12]

$$\textit{occupancy}_{it} = \alpha_0 + \alpha_1 * \textit{price}_{it} + \alpha_2 * \textit{rating}_{it\text{-}1} + \alpha_3 * \textit{average\_occupancy}_{it\text{-}1} + \\ \alpha_4 * \textit{post\_switch}_t + \alpha_5 * \textit{inst\_book\_group}_i * \textit{post\_switch}_t + \alpha_6 * \textit{host\_controls}_i \quad (2.2) \\ + \alpha_7 * \textit{listing\_controls}_{it} + \alpha_8 * \textit{month\_dummy}_t + w_i + \varepsilon_{it}$$

$$\textit{price}_{it} = \beta_0 + \beta_1 * \textit{average\_occupancy}_{it\text{-}1} + \beta_2 * \textit{rating}_{it\text{-}1} + \beta_3 * \textit{post\_switch}_t + \\ \beta_4 * \textit{inst\_book\_group}_i * \textit{post\_switch}_t + \beta_5 * \textit{host\_controls}_i + \beta_6 * \quad (2.3) \\ \textit{listing\_controls}_{it} + \beta_7 * \textit{month\_dummy}_t + w_i + \eta_{it}$$

$$\textit{rating}_{it} = \gamma_0 + \gamma_1 * \textit{price}_{it} + \gamma_2 * \textit{occupancy}_{it} + \gamma_3 * \textit{post\_switch}_t + \\ \gamma_4 * \textit{inst\_book\_group}_i * \textit{post\_switch}_t + \gamma_5 * \textit{rating}_{it\text{-}1} + \gamma_6 * \textit{host\_controls}_i + \gamma_7 * \quad (2.4) \\ \textit{listing\_controls}_{it} + \gamma_8 * \textit{month\_dummy}_t + w_i + \theta_{it}$$

In the above model, $\textit{inst\_book\_group}_i$ is a dummy variable indicating if listing i is in the treated group ($\textit{inst\_book\_group}_i=1$) or in the control group ($\textit{inst\_book\_group}_i=0$). The dummy variable $\textit{post\_switch}_t$ denotes post-switching period vs. pre-switching period ($\textit{post\_switch}_t=1$ if $t>0$ and 0 if $t<0$ for every treated listing and its matched control listing). The term w_i captures time-invariant listing fixed effects and $\textit{month\_dummy}_t$ captures seasonality effects.[13] Equation (2.2) specifies the occupancy of listing i in month t is affected by its price in the same month, review rating score and average occupancy before month t. Because a listing's price is

---

[12] Of the 369 treated samples that are not matched, 144 listings are removed for not being on common support and another 225 listings are dropped due to failure of getting a match.

[13] $\textit{month\_dummy}_t$ captures fixed effects related to month (e.g., January vs. February vs. March). While January may be a pre-treatment month for some listings that switched later, it may be a post-treatment month for other listings that switched earlier. Therefore, introducing $\textit{month\_dummy}_t$ does not lead to dropping $\textit{post\_switch}_t$ from the model.

typically set up by the host on a month-by-month basis in advance, I specify that the price of listing i in month t is affected by its lagged average occupancy and rating score in equation (2.3). Finally, considering that $rating_{it}$ measures the cumulative review rating of listing i by the end of month t, in equation (2.4) I model it as a function of its starting value before month t ($rating_{it-1}$) as well as other factors such as the listing's price and occupancy in month t that could affect newly obtained ratings. In addition, I control for several time-varying variables that could affect a listing's performance. For example, a listing with more cumulative number of reviews by the previous month ($review\_count_{it-1}$) may be perceived as more reliable by guests, leading to higher occupancy and the host's ability to charge a higher price in the current month. But meanwhile, the associated higher expectations from the guests may cause the rating of the listing to drop down in the current month. The major coefficients of interest are α_4, β_4, and γ_4, which measure the causal impact of forgoing screening on a listing's occupancy, price, and review rating respectively.

Second, a key assumption of the DiD estimation is that the control and treated groups have a similar time trend in the absence of the treatment. To test the parallel trend assumption, I follow prior research (e.g., Autor 2003; Ryan et al. 2015) and model the time trends of a listing's performance by introducing interactions of the treatment indicator and various time dummies as shown below:

$occupancy_{it} = \alpha_0 + \alpha_1*price_{it} + \alpha_2*rating_{it-1} + \alpha_3*average\_occupancy_{it-1}$
$+\alpha_4*host\_controls_i + \alpha_5*listing\_controls_{it} + \alpha_6*month\_dummy_t +$ $\qquad$ (2.5)
$\sum_j \alpha_j^p *inst\_book\_group_i*relative\_month_j + w_i + \varepsilon_{it}$

$price_{it} = \beta_0 + \beta_1*average\_occupancy_{it-1} + \beta_2*rating_{it-1} + \beta_3*host\_controls_i +$
$\beta_4* listing\_controls_{it} + \beta_5* month\_dummy_t +$ $\qquad$ (2.6)
$\sum_j \beta_j^p *inst\_book\_group_i*relative\_month_j + w_i + \eta_{it}$

$$rating_{it} = \gamma_0 + \gamma_1*price_{it} + \gamma_2*occupancy_{it} + \gamma_3*rating_{it-1} + \gamma_4*host\_controls_i$$
$$+ \gamma_5*listing\_controls_{it} + \gamma_6*month\_dummy_t + \qquad\qquad (2.7)$$
$$\sum \gamma_j^p*inst\_book\_group_i*relative\_month_j + w_i + \theta_{it}$$

where, *relative_month_j* is a dummy for each month *j* relative to the treatment month 0. Appendix Table A4 presents the estimation results. An F-test fails to reject the null hypothesis that all the pre-treatment coefficients are jointly zero, providing support for the parallel trend assumption.

Finally, a third challenge to my estimation is simultaneity in the determination of occupancy, price and review ratings. Given that endogenous dependent variables such as *occupancy_{it}* and *price_{it}* also enter the estimation equations on the right hand side and the three error terms may be correlated with each other due to unobservable common factors related to a listing, I simultaneously estimate the three equations by employing the three-stage least squares (3SLS) approach which has been commonly used in estimating a system of equations in similar situations (e.g., Aral et al. 2018; Chang and Gurbaxani 2013; Roberts and Schlenkera 2013; Xue et al. 2012). In economic theory, it is reasonable to assume that competing products' characteristics have no direct impact on a consumer's utility for a product but may directly affect the characteristics of the product, which are endogenously decided based on competitions (Berry et al. 1995). Similarly, Hausman (1996) argues that the variations in the equivalent characteristics of competitors or neighbors should correlate with the focal player's characteristics but would only affect other characteristics of the focal player indirectly. Echoing these studies, Reynaert and Verboven (2014) discuss and establish why competitors' characteristics are valid instruments for the focal player's price. Therefore, a variety of studies have used the characteristics of competing products as

instrument variables, sometimes referred to as *Hausman type instruments*, for a product's endogenous characteristics (e.g., Aral et al. 2018; Cachon et al. 2019; Fan 2013; Ghose and Han 2014; Hong and Pavlou 2016; Xue et al. 2012). Following the same approach, I use the average price of a listing's 12 local competitors as an instrument for *price*$_{it}$. Similarly, I use the average occupancy of a listing's closest local competitors as an instrument for *occupancy*$_{it}$.[14]

2.4.2.2 Estimation Results

The outcomes of the DiD estimation are presented in columns (1) to (3) of Table 5. In terms of control variables, I observe positive and significant coefficients of *review_count*$_{it-1}$ in the occupancy and price equations. However, *review_count*$_{it-1}$ has a negative and significant influence on a listing's rating. In addition, the coefficient of *host_is_professional*$_{it}$ is positive and significant in the price equation, indicating that professional hosts that are experienced in the hospitality business such as hotels generally charge higher price than non-professional hosts do.

More importantly, I find that forgoing screening significantly improves the monthly occupancy level of a listing by 13.52% (α_4=0.073, 0.073/0.54=13.52%) or about 1.79 nights (0.073*24.5=1.79), as indicated by the positive and significant coefficient of the interaction term *inst_book_group*$_i$**post_switch*$_t$. Such an increase in occupancy improves the revenue of the host by an average of $300.23

---

[14] The Kleibergen-Paap rk LM statistic is 93.26 (p<0.001) in the occupancy equation and 143.53 (p<0.001) in the rating equation, providing evidence for instrument relevance. The Cragg-Donald Wald F statistic is 13541.68 in the occupancy equation and 60.96 in the rating equation, exceeding the recommended Stock-Yogo critical value for 10% maximal IV size in each equation and thus providing support that my instruments are not weak. In unreported results, we also use the lagged value of the endogenous variables as excluded instrument variables (e.g., Xue et al. 2012; Chang and Gurbaxani 2013; Fuentelsaz et al. 2012) and obtain consistent estimation results.

(1.7885*$167.87=$300.23) per month. Second, I find that the price of a listing does not change after the host forgoes screening. However, the review rating of the listing falls by 1.078 points or 1.20% (-1.078/89.55=-1.20%) after forgoing screening, as indicated by the negative and significant coefficient of the interaction term $inst\_book\_group_i*post\_switch_t$.

Figure 2.2 illustrates the causal impacts of forgoing screening based on the DiD estimation results: compared to the matched control listings, the treated listings have a significant increase in the average predicted occupancy, negligible change in the average predicted listing price, and a salient decrease in the average predicted rating.

To investigate whether the impacts of forgoing screening depend on host race/ethnicity or gender, I add the three-way interaction terms $inst\_book\_group_i*post\_switch_t*race/ethnicity_i$ and $inst\_book\_group_i*post\_switch_t*gender_i$ to the DiD estimation model. The baseline comparison groups are hosts whose race/ethnicity are unidentified in terms of host race/ethnicity, and female hosts in terms of host gender, respectively. The main results with both moderating effects included is presented in Table 2.6. The coefficients of the three-way interaction terms for African American hosts and for White hosts are insignificant in both occupancy and price equations. Considering the lower monthly occupancy of African American hosts before they forgo screening (see Appendix Table A2), African American hosts have a higher percentage of occupancy improvement (0.079/0.43=18.37%) than White hosts (0.079/0.54=14.62%). In the equation for review rating, the coefficient of the three-way interaction term is significant with a magnitude of 1.149 for African American hosts but is insignificant for White hosts. These results indicate that African American

35

hosts incur a lower cost to reap the higher percentage of occupancy improvement—on average, they experience a 1.149 points or 1.28% lower decline in review ratings (0.28% decline for African American hosts versus 1.56% decline for White host). Turning to the moderating effects of host gender in Table 6, I focus on the three-way interaction terms involving host gender, where the baseline comparison group is female hosts. The coefficient of the three-way interaction term for male hosts is only significant in the price equation with a magnitude of -4.280 (p=0.001). This indicates the listing price of a male host decreases by $4.28 compared to a female host; but there are no significant differences in either occupancy rates or review ratings across gender. In unreported results, I obtain consistent results when examining the moderating role of host race/ethnicity and host gender independently.

Figure 2.3 highlights the outcome heterogeneity on top of the average effects of forgoing screening based on the DiD estimation results. It shows that: the average predicted price of a male host's listing decreases as compared to that of a female host's listing; and the average predicted listing rating drops less for African American hosts than for White hosts.

2.4.2.3 Uncovering What Drives Outcome Heterogeneity

What drives such heterogeneity in outcomes? I examine if the differential benefits and costs of forgoing screening for White hosts as compared to African American hosts can be explained by their differences in quality as reflected in review ratings and comments.[15] As shown in Table 2.7, among switchers, African Americans are not different from White hosts despite that White hosts in general have higher

---

[15] We focus on listings that have comments in both pre-switching and post-switching periods and have a minimum of 5 comments so that we could reliably analyze their textual comments.

review ratings. This suggests that White switchers are of lower relative quality (vis-à-vis an average White hosts) than African American switchers (vis-à-vis an average African American hosts). I also analyze the textual review comments received by hosts. For each comment, I identify the sentences with positive sentiment as positive mentions, and sentences with negative sentiment as complaints that are further classified into various categories [16] (e.g., complaints about room, complaints about noise, complaints about location, etc.). As shown in Table 7, there are no significant differences between African American hosts and White hosts in positive mentions per comment and negative complaints per comment before they forgo screening. After switching to the forging screening state, both African American and White hosts suffer from reduced positive mentions in the review comments. However, African American hosts have higher positive mentions per comment than White hosts while their differences in complaints remain insignificant. These findings are consistent with my DiD estimation result that the cost of decreased ratings for African American hosts is lower. Overall, the analyses on review ratings and comments provide evidence that the quality of African American switchers (relative to an average African American host) is higher than White switchers (relative to an average White host). This quality difference enables African American hosts to reap more benefits from forgoing screening at lower costs than White hosts.

---

[16] For this, we first select 1600 comments left by guests, while ensuring that host gender and race/ethnicity are balanced. We then manually tag the 9925 individual statements from these comments as either positive mentions or complaints. Furthermore, we categorize the identified complaints into 6 major types. We then train a LSTM (long short-term memory) based Recurrent Neural Network to classify individual statements. A 10-fold cross validation gives over 90% accuracy of my classification. We then employ this trained model to detect complaints and positive mentions in over 491,000 comments left by guests for the hosts in my sample.

To further understand why female hosts and male hosts are similarly affected by forgoing screening except for differences in listing price change, I compare their quality and listing prices before forgoing screening. The results are shown in Table 2.8. I find that overall female and male hosts are of similar quality; however, female hosts that forgo screening have, on average, a higher rating than male hosts that forgo screening. A comparison of the listing prices of female switchers and male switchers at various occupancy ranges shows that female hosts have consistently lower listing prices, albeit having better ratings, than male hosts. This result echoes prior studies on gender gap in online markets, which suggest that female sellers suffer from lower prices than male sellers (e.g., Kricheli-Katz and Regev 2016). Taken together, the results suggest that due to their lower quality than female switchers, male switchers would have smaller improvement in occupancy and greater decrease in ratings after forgoing screening. However, this is offset by male switchers' act of lowering listing prices after forgoing screening. As a result, I observe similar changes in occupancy and ratings across the two genders but a drop in listing price for male switchers as compared to female switchers.

2.4.3   Falsification Tests

Overall, my analyses suggest that: (1) one major reason for hosts to voluntarily forgo screening is to improve listing occupancy, albeit at the cost of lower review ratings; and (2) White hosts incur higher costs and improve the occupancy levels of their listings less than African American hosts after forgoing screening. Following the same cost-benefit calculus logic, these results imply that listings with higher costs of forgoing screening and listings by White hosts are more likely to switch back to the screening state. Therefore, as a falsification test, I also examine which listings among switchers

38

are more likely to switch back to traditional screening (which I call *two-way switchers*), relative to continuing with forgoing screening. The sample for this analysis includes observations for all switchers after they have forgone screening. To account for the fact that the value of the dependent variable (*traditional_bookable$_{it}$*) will always be 1 after a two-way switcher has switched back to traditional booking, I also only keep two-way switchers' observations in the month of switching back while dropping their observations afterwards. Formally, I model the decision to switch back a listing to traditional booking as:

$$\textbf{\textit{Logit(traditional\_bookable}}_{it}\textbf{\textit{)}} = \pi_0 + \pi_1 * occupancy\_change\_percent_{it-1} + \pi_2 * price\_change\_percent_{it-1} + \pi_3 * rating\_change\_percent_{it-1} + \pi_4 * host\_controls_{it} + \pi_5 * listing\_controls_{it} + \pi_6 * competitor\_controls_{it} + \pi_7 * month\_dummy_t + \varepsilon_{it}$$

(2.8)

Similar to equation (2.1), here I model a listing's decision to switch back to traditional booking in month *t* as a function of variables measuring how much a listing has benefited from forgoing screening, including change in occupancy (*occupancy_change_percent$_{it-1}$*), change in listing price (*price_change_percent$_{it-1}$*), and change in cumulative review rating (*rating_change_percent$_{it-1}$*), and other host-related and listing-related variables.

I report the estimation results in Appendix Table A5. I observe insignificant coefficients of *occupancy_change_percent$_{it-1}$* and *price_change_percent$_{it-1}$* in column (1). However, I find that the coefficient of *rating_change_percent$_{it-1}$* is significant and negative, suggesting that listings that experience greater drops in rating due to forgoing screening are more likely to re-enable screening. Overall, the results suggest no discernable differences effects of occupancy and price changes, but listings that experience a higher drop in review ratings and thus incur a higher cost from forgoing

screening are more likely to switch back to screening.

African American hosts, I find, have a lower propensity to re-enable screening as compared to White hosts. This result, along with finding on *rating_change_percent*$_{it-1}$ (that listings with higher drop in ratings are more likely to re-enable screening), confirms my conjecture and provides further support for the main findings regarding the antecedents and consequences of forgoing screening. I find no significant difference between female hosts and male hosts in their propensities to re-enable screening. However, in additional sub-sample analyses (see columns (2) and (3) of Appendix Table A5) I find that among females, African American females are more likely to continue to forgo screening than White females. Similarly, among males, African American males are more likely to stay with forgoing screening than White males.

2.5   Discussion

2.5.1   Key Findings and Discussion

Motivated by the growing implementation of the forgoing screening mechanism in online P2P markets and a dearth of research examining when and who is likely to do so, and what are the resultant outcomes, this study offers a concerted and systematic empirical investigation into these questions in the context of Airbnb, a popular P2P sharing platform. The findings are summarized in Table 2.9. I find demand conditions facing individual listings are associated with this choice: when listings experience mid-level or declining occupancies, the expected benefits of forgoing screening seem to outweigh the expected costs in the decision-making calculus and such listings are more likely to forgo screening. This is also true for listings with higher local competition, and for listings where the local competitors have forgone screening, suggesting that listings may forgo screening due to competitive pressure. Host characteristics also

matter. Supply characteristics of hosts affect the choice: new entrants on the Airbnb platform and non-superhosts are also more likely to forgo screening. These results suggest that inexperienced hosts may leverage instant booking to learn evaluating the quality of guests. I also find that professional hosts are more likely to forgo screening, probably because those professional hosts are used to and comfortable with the practice of not screening their guests. Importantly, and consistent with the literature on discrimination and bias, I find non-economic characteristics of hosts such as race/ethnicity and gender are associated with forgoing of screening: African American hosts, and female hosts are more likely to forgo screening.

The outcomes associated with a choice to forgo screening are consistent with the expected costs and benefits. Occupancy rates indeed increase for listings with no associated changes in price, thus increasing the total revenue for the listing. However, as would be expected given that providers open the possibility of renting to unknown and potentially low-quality guests, or to less well-matched guests, the reviewer ratings also decline after forgoing screening. Importantly, I find African American hosts benefit in terms of greater percentage increases in occupancy rates, and also pay lower penalty in terms of smaller declines in review ratings, relative to White hosts. This is mainly because African American switchers are of higher quality than White switchers among their same-race pool of hosts. In addition, while female hosts are no different from male hosts in terms of changes in occupancy rates or review ratings, they do not suffer from the cost of decreased listing price as male hosts who are of worse quality. Thus, the benefits of forgoing screening accrue disproportionally more for African American and/or female hosts. Given guests are most likely aware of hosts'

ethnicity/gender based on their profile photos, this suggests that guests seem to tradeoff their taste-based bias for the added convenience of instant booking.

2.5.2   Limitations and Future Research

I acknowledge several limitations of this study. First, similar to prior research on Airbnb (e.g., Edelman and Luca 2014), my sample consists of all Airbnb listings in New York City within a specific timeframe (i.e., from August 2015 to February 2017). While New York City represents the largest locale in terms of active US Airbnb listings, surpassing the next biggest city Los Angeles by 55% more rental supply, generalizability concerns limit applying the findings of this study to other cities, particularly international locations where there may be significant differences in socio-economic characteristics. We need additional studies across locations to verify consistent findings, or identify contingency factors affecting the relationships I document. In particular, future studies could compare New York City with another city for which the "Instant Book" feature was popularized at a different time to further explore how the availability or adoption of forgoing screening exerts influences. Similar concerns apply when generalizing findings to other online P2P platforms. In my context, providers could switch relatively easily. Caution should be exercised to ensure online P2P platforms conform to the boundary conditions that the markets represent relatively low switching costs between implementing or forgoing screening. For platforms with high cost of switching or lack of autonomy to switch for participants, the cost-benefit calculus of participants may be more complicated and thus further studies may be needed to verify the applicability of my findings.

Second, while I focused on heterogeneity in listing and host characteristics, my study abstracted away from demand heterogeneity, and potential market segmentation

with sub-pools of buyers and suppliers who equilibrate around choices of implementing or forgoing screening. Data constraints preclude the examination of whether different screening choices represent shifting compositions in guests, such that different guests find non-screening desirable. Particularly worthy of attention is whether instant booking enables greater substitution of Airbnb to traditional hotels where guests are not constrained by provider screening when making their choices. Future studies could compile data on buyers on Airbnb to append to the data on providers, and examine how differences in composition or potential increases in demand relate to provider choices of implementing or forgoing screening. Such data would also enable future research to shed light on the impact of one party forgoing screening from the other party's perspective, or the impact on market efficiencies in implementing options for forgoing screening for one or both sides of the market. Here, studies may be able to examine how buyers, and other providers, strategically respond to the focal provider's choice of screening. For example, while I took a focal provider (listing)'s perspective on the option of forgoing screening, it is possible that the guests on Airbnb will have strategically changed their behavior too. Prior research suggests participants' game with mechanisms in online markets (e.g., Klein et al. 2009; Ye et al. 2014), and hosts' forgoing screening may induce the guests to behave more opportunistically.

Third, one potential reason for Airbnb's introduction of the forgoing screening mechanism is to combat observed accounts of racial discrimination.[17] While forgoing screening clearly precludes hosts from engaging in racial discrimination, my study found evidence that race/ethnicity, more so than gender, matters in both whether hosts

---

[17] https://www.airbnbcitizen.com/a-fair-community-for-everyone/

choose to implement or forgo screening, and the associated outcomes of their choices. The lack of granular data on interactions between hosts and guests precludes me from examining the mechanisms causing these choices and outcomes, and a deeper investigation into these issues represent important future research avenues. For example, it is not clear whether the results that African American hosts are more likely to forgo screening is because they are less risk-averse as compared to Whites (e.g., Gutter and Hatcher 2008; Halek and Eisenhauer 2001), or whether, as I assume, they are responding to underlying heterogeneity in their individual demand, holding constant risk preferences. Similarly, differential outcomes associated with choices when exercised by hosts of different race/ethnicities may be reflective of statistical or taste-based discrimination, and my study was unable to discern which factor dominated. As an example, it is unclear whether the lower declines in review ratings for African American hosts could also be driven by the guests' higher ex ante expectations for Whites than for African Americans (Doleac and Stein 2013), causing higher possibility of ex post disconfirmation from unscreened guests and thus greater damage in review ratings for White hosts. Future studies could utilize laboratory or field experimental design to isolate and identify if other underlying mechanisms are at play for the results I uncover. Such studies would be critical for a deeper understanding of how market design on online platforms impact social welfare. For example, currently Airbnb automatically prioritizes instantly bookable listings in the search results. Could this lead to a disadvantage for White hosts who gain less from forgoing screening? Such questions highlight the importance of performing a comprehensive evaluation of any market mechanism by examining its direct and extended impacts for participants

44

of different races/ethnicities and genders. A careful identification of tradeoffs is needed to ensure overall benefits outweigh potential side effects and costs, as they relate not only to economic, but also social consequences in the form of discrimination based on race/ethnicity or gender.

Fourth, while I have examined what listings are more likely to forgo screening and excluded several alternative explanations of why a listing may forgo screening, more studies leveraging causal designs can help clearly identify the mechanisms that drive a listing's decision. Finally, while I have used different matching methods in the DiD estimation to account for the endogeneity issue of a listing's forgoing screening status in a quasi-experimental setting, it is plausible that the self-selected status of a listing is influenced by other unobservable factors. Future research could leverage exogenous shocks on the Airbnb platform to further examine a listing's decision of forgoing screening and its impacts in a natural experiment setting.

Chapter 3: Delaying Informed Consent: An Empirical Investigation of Mobile Apps' Upgrade Decisions

3.1   Introduction

The past decade has witnessed a marked rise in consumer sensitivity to online information collection and privacy. About 86% of internet users have taken steps to avoid surveillance by organizations during their online browsing sessions (Rainie et al. 2013). Recent cases of massive breach in privacy such as the Cambridge Analytica scandal (Cadwalladr and Graham-Harrison 2018) have only amplified such user concerns. With the rapid proliferation of mobile devices, privacy concerns have naturally extended to mobile devices (Barkhuus and Dey 2003). Smartphone mobile apps have traditionally obtained blanket permissions from users to access their sensitive information upon downloading these apps. IT security researchers have shown that over a third of the apps over-seek permissions (i.e., seek more permissions than strictly required for their functionality), which increases the risk of data misuse (Felt et al. 2011). Users have increasingly become sensitive to such practices and have been proactively taking measures to protect their privacy, including abandoning apps or abandoning the platform altogether (Pingitore et al. 2017).

In an attempt to respond to these concerns and provide users with better control over their mobile footprint, mobile platforms such as Android have released upgrades that provide users with fine-grained control over their information. In this study, I examine one such upgrade – the release of Android's version 6.0 in late 2015. In earlier versions of Android, upon downloading the app, users automatically agreed to provide access to all sensitive information listed upfront by the app (such as contacts, phone memory, GPS). Android 6.0 allowed users to download apps without granting access

to any sensitive information and then required the apps to seek individual permissions to access users' sensitive information (viz. dangerous permissions[18], hereinafter simply referred to as permissions[19] for the sake of brevity) when users opened the app. Such a change in Android's privacy policy (illustrated in Figure 3.1) provides users with a choice to use the "watered-down"[20] version of the app by granting permissions to sensitive information "as they desire".

While such changes to security and privacy policies are not new among platforms, Android's decision to provide a time-window for apps to adhere to the new policy, makes this policy change unique and worth investigating. Before Android, Apple had introduced a fine-grained permission structure for sensitive information in iOS in 2012 while Microsoft had introduced User Access Control in Windows Vista OS in 2007 where programs required explicit permissions from users to access OS kernel. In both these instances, platforms used their market power to force all apps to upgrade and adhere to such policy changes to be eligible to run in their latest version. In contrast, Android gave mobile apps a window of three years (anytime between 2015 and 2018) to upgrade, instead of requiring them to upgrade immediately to the latest version[21]. While iOS and Windows required all apps to target the respective latest versions to be

---

[18] As per Android's website, "Dangerous permissions cover areas where the app wants data or resources that involve the user's private information, or could potentially affect the user's stored data or the operation of other apps"
https://developer.android.com/guide/topics/permissions/overview.html#normal-dangerous
[19] Apps are only required to explicitly seek dangerous permissions while they do not have to seek normal permission (e.g., phone vibration, prevent screen lock, turn on camera flash etc.).
[20] Users are not required to grant all permissions just to download the app. Instead, users can grant partial or no permissions but still download and run the app, to the extent that such permissions are not needed for its functioning.
[21] At the outset, we differentiate upgrades by users and apps. Users' platform upgrade means that they install the latest version of Android on their mobile device. App's platform upgrade means that apps compile using the latest version of the Software Development Kits (SDKs) provided by the platform and choose the latest version of the platform they want to *target* (called *targetAPI* in Android).

eligible to run on upgraded devices, Android allowed non-upgraded apps (i.e., targeting a lower Android version) to run on the latest Android phones.

Android apps' choice of when to upgrade to 6.0 (anytime between 2015 and 2018) meant that only those apps that upgraded to the latest version were required to adhere to the new privacy policies. Apps that delayed upgrading, however, could continue to seek permissions at download (instead of at run-time) even when such apps ran on phones that were upgraded to the latest version. Such apps could also release feature updates or fix bugs using older software development kits (SDKs). In other words, apps that delay upgrade would be "fully functional". However, the apps that stayed with the older version of Android would forgo access to the latest platform features, optimizations, and support.

Given the tradeoff between upgrading early to benefit from the platform's latest features and staying with the older version to retain blanket access to user information, my study examines the choices made by different types of apps and their consequent outcomes. I first perform an exploratory analysis to understand the driving factors that prompt apps to delay upgrade. Next, I perform casual analysis to identify the effects of delaying upgrade on apps' marketplace outcomes, i.e., their popularity (that dictates their visibility on the marketplace) and the user rating. Specifically, the three broad research questions that this project addresses are: (*i*) *What are the characteristics of apps that delay upgrading to seeking run-time permissions? (ii) Does the (non)essentiality of the certain permissions to app's working impact app's decision to delay upgrade?,* and *(iii) How does strategically delaying upgrading affect apps' marketplace outcomes (app rating and app popularity)?.*

I assemble a unique panel of apps between April 2016 and March 2018 by installing 13,604 most popular free Android apps on emulators and updating these apps on a monthly basis. I specifically develop an android app that scans the android app database of the host emulator and determines the Android version that each app targeted, along with a list of permissions sought by the apps. Next, I collect monthly snapshot of app characteristics, such as download buckets, ratings, total comments, categories, file-size (in MB), screenshots, app description, revenue model and developers' information of over 2 million Android apps.

To examine which apps over-seek permissions to users' sensitive information, I categorize permissions into two groups: permissions that are required for apps' operations (*essential* permissions) and those that are not (*non-essential* permissions). I employ Skip-gram Word2Vec and k-means clustering to divide apps from each Android app category into sub-categories based on their functional similarity and then statistically determine if each permission sought by apps is essential or not. Further, I employ weighted PageRank algorithm to derive a measure of apps' popularity. For generating the PageRank, I use apps' "similar apps" list on Play Store.

To preview my findings, I employ Cox Proportional Hazard models and find apps that over seek permissions are more likely to delay upgrading to Android 6.0. Further, in investigating how the essentiality of each of the permissions for each app affects their propensity to delay upgrade, I find that the likelihood of delaying upgrade increases only for apps that seek more contextually *non-essential* permissions, while seeking *essential* permissions does not increase the propensity to delay upgrades. In examining further to understand what kinds of apps are more likely to display over-

seeking behavior, I find that apps with in-app advertising as a revenue model are more likely to over seek *non-essential* permissions. Here I complement econometric methods with modern machine learning techniques to ascertain the apps' over-seeking behavior. In essence, I find that apps that benefit from the use of sensitive information for in-app advertising revenues are more likely to continue over-seeking non-essential permissions and delay upgrading to Android 6.0. To address the concern that app's decision to delay its upgrade could be operational, rather than strategic, I (a) employ models with developer fixed effects, (b) conduct a falsification test, (c) employ alternative (i.e., logit) model specifications and (d) forward evidences to rule out demand side reasons, supply side cost-reasons or lack of benefits reasons to delay upgrade.

By employing popular matching techniques such as Dynamic Propensity Score Matching along with Difference-in-Differences technique, I find that apps that delay upgrading to the latest version significantly reduce seeking *non-essential* permissions by 6.79%, a larger magnitude reduction in seeking *non-essential* permissions than on-time upgraders among apps. I also find that delaying upgrades effectively erases potential gains in ratings that apps would have received after upgrading, an equivalent of losing a rating boost from about 560 users. Furthermore, delaying upgrading to the latest Android version significantly reduces apps' popularity measure by 9.06%, suggesting that strategic delays in upgrading to Android 6.0 results in both monetary (lower popularity score leading to lower downloads) and reputational (lower rating) repercussions.

This study makes a number of significant contributions. While extant research has investigated buyers' perspective to changes in information privacy policies, to the best of my knowledge, this is the first paper to investigate sellers' responses to exogenous changes in information gathering practices. Next, I add to the literature on temporal choices by showing that upgrading to the latest version at the earliest leads to beneficial outcomes for the apps in the marketplace, despite potentially disrupting its revenues in the short run. Next, my study guides platform mechanism designs by showing that providing a meaningful choice to app users to control their sensitive information improves apps' information collection behavior. The findings suggest that reducing the cost of making privacy choices may help address the privacy paradox documented in privacy literature. By showing that even popular apps systematically lower seeking permissions, I highlight the role of platform mechanisms in safeguarding consumer interest. In doing so, I introduce the technique of utilizing advanced text analysis techniques to identify peer apps. Finally, my research adds to the literature on platform strategy by suggesting that a fragmented platform such as Android should carefully design the upgrade window to mitigate sub-optimal outcomes for the participants (apps, users) or the platform.

3.2   <u>**Theoretical Background**</u>

This research draws upon and contributes to three distinct streams of literature: users' responses to privacy policies, user privacy preferences and strategic timing by firms. To place the contributions in perspective, I review the relevant literature, identify the gaps and define my research questions.

3.2.1 Users' responses to Privacy Policies

Consumers' responses (or lack thereof) to privacy concerns in an online market have been extensively studied in literature (see Bélanger and Crossler (2011) for a survey of privacy research in IS field). Consumers' concerns about information privacy may stem from their limited control on how their personal information may be used by the online seller (Acquisti and Grossklags 2005b, Dinev and Hart 2006, Pavlou et al. 2007). Such concerns are quite valid, especially when the secondary usage of their private information may come at a cost to them in their future transactions (Varian 2009). Privacy concerns may also arise when consumers perceive sellers' data gathering activity as excessive or irrelevant. Such perceptions impact consumer beliefs about trustworthiness of sellers and risks related to interactions with such sellers (in this case, apps) (Chellappa and Sin 2005, Malhotra et al. 2004, Pavlou et al. 2007).

Under conditions that induce privacy concerns, users have responded by expressing preference to and/or taking actions to control of their sensitive information. One sub-stream of literature has investigated consumers' privacy concerns independent of sellers' actions. Researchers have documented surveys and observations where significant proportion of consumers seek to gain control over information gathering and secondary usage of such information (Acquisti and Grossklags 2003, Hoffman et al. 1999, Phelps et al. 2000). Another sub-stream has investigated users' reactions when sellers choose to adopt privacy practices. For instance, prior work (Aguirre et al. 2015, Culnan and Armstrong 1999, Dinev and Hart 2006, Hui et al. 2007) has shown that users' risk perceptions and privacy concerns are significantly alleviated when websites guarantee fair information practices. Allowing consumers to opt-in for data collection also significantly increase their responses to personalized advertisements (Nowak and

52

Phelps 1995, Tucker 2014) or increase utilizing personalized services (Aguirre et al. 2015, Chellappa and Sin 2005, Nowak and Phelps 1995, Tucker 2014).

While there is a substantial body of research on users' responses to privacy changes, hardly any studies have examined sellers' responses to changes in privacy policies when such policies are exogenously imposed upon them. Literature has documented sellers' response to changes in the nature of competition in the market (Chen and Forman 2006, Foerderer et al. 2018), however these responses are voluntary. Recent studies have investigated firms' responses to adhering to the General Data Protection Regulation (GDPR) compliance (Garber 2018, Gradwohl 2018) but none has studied sellers' responses in the context of online or mobile platforms. This research contribute to this stream of literature by examining apps' responses when the Android platform enforces a new privacy policy. I document robust evidence that apps with questionable data collection behavior exhibit a tendency to strategically delay an upgrade that limits such behavior.

3.2.2 Users' Privacy Preferences

Despite an increase in users' concerns about their privacy, studies have consistently documented a paradox when such consumers internalize privacy choices. Researchers (Acquisti and Grossklags 2003, 2005b, Hann et al. 2007, Spiekermann et al. 2001) have observed that users that express concerns about privacy do not act to address such concerns when they have a choice. Some of the common inactions include not reading the privacy policies, not going through privacy settings, or not understanding the implications of Android permissions (Hoofnagle et al. 2010, Jensen and Potts 2004, Kelley et al. 2012). Furthermore, studies have found that buyers that express desire to protect their privacy are willing to exchange/disclose sensitive information for short

term benefits (Acquisti and Grossklags 2005a, Hui et al. 2007). Also, despite stating their desire to pay a premium for privacy (Tsai et al 2011), buyers have demonstrated a lack of willingness to incur a cost for privacy when such an option is made available to them (Brunk 2002, Rose 2005). Such inconsistencies in internalizing privacy choices by users may be explained by the difficulty in gauging the trade-off between the cost incurred from disclosing their sensitive information versus the benefits derived from disclosing sensitive information (Acquisti and Grossklags 2005b).

Despite a large volume of existing literature on users' privacy behavior, not many studies have examined the changes in users' privacy behaviors when the cost of making the above described tradeoff is reduced. This research add to this literature stream in two ways. I show that reducing users' cost of making privacy choices helps improve sellers' information gathering practices on platforms. I show that shifting to (run-time) *à la carte* permission seeking reduces users' cost of making privacy choices by allowing users to evaluate the relevance of each permission before granting them. I also find that apps reduce seeking permissions to access sensitive information after the Android update.

3.2.3 Early vs. Late Movers

My study is also related to a vast body of research examining firms' timing of entry into new markets. Researchers (Barnett et al. 2012, Kalyanaram and Urban 1992, Kerin et al. 1992, Mitchell 1991) have studied early mover advantages for firms and have shown that early movers benefit from sustained market-share advantages. Other researchers (Agarwal and Gort 2001, Carow et al. 2004, Lieberman and Montgomery 1988) have found that under certain conditions such as low imitation costs (absence of resource asymmetries, weak IP regime etc.), high market uncertainty (lack of industry

standard, large technological discontinuities etc.) or rapid interfirm diffusion of technology (high labor mobility, improvement in communication etc.), late movers might benefit more. Most of these studies examine firms' timing of entry in a market setting. In contrast, my study focuses on the timing of policy adoption, rather than market entry. While market entry decisions are often made to maximize profits and/or gain market share, decisions regarding the timing of policy adoption are often made to minimize potential losses due to policy enactments. Hardly any studies have examined firms' temporal choices under a policy enactment setup. Recent studies that examine firms' responses to governmental policies such as GDPR have shown that well-informed market players may be able to utilize the timing of movement as a signal of trustworthiness and a differentiator (Garber 2018, Gradwohl 2018). I am unaware of similar studies on an online platform setting. In essence, literature on early movers/adopters of change points towards a lack of rigorous investigations studying the timing of policy adherence on medium- to long-term revenue and reputational implications on the platform. I contribute to this literature by investigating the timing of policy adoption decisions by apps in a mobile platform, and their impacts on outcomes.

3.2.4 Research Questions

I find several gaps in literature that reference understanding the antecedents and consequences of delaying privacy policy adoption on platforms. First, there is a dearth of research studying factors that drive some of the apps' upgrade decisions. Hence, my first research question (RQ1) asks and answers, *"What are the characteristics of apps that delay upgrading to seeking run-time permissions?"* I premise this question using the insights from literature on responses to privacy features. The shift from blanket

permissions at download-time to *à la carte* permissions at run-time focuses users' attention on individual permissions sought by apps. Such enhanced attention on permissions may drive apps that seek a lot of permissions to strategically delay upgrading to the latest version.

In my second research question (RQ2), I further test my conjecture about the strategic behavior of apps by distinguishing whether each of the permissions sought by apps are contextually essential (for app's working) or not. As an illustration, in the context of a ride-hailing app (such as Lyft/Uber), accessing GPS is contextually essential but accessing phone's microphone may be contextually irrelevant. Specifically, I ask and answer, *"Do strategic reasons (such as monetization from continued access to non-essential (to app's core functionality) sensitive information) drive app's decision to delay upgrading to seeking run-time permissions?".* I premise this question by referring to the information security literature that has documented permission over-seeking behavior. While upgrading to the latest version of Android provides apps with access to the platform's latest features and improvements, it also restricts free access to users' sensitive information that apps enjoyed prior to upgrade. It makes sense then, that upgrading to the latest version has different associated costs and benefits for apps with different revenue models. Based on this cost-benefit calculus logic, I conjecture that apps that benefit from over-seeking *non-essential* permissions from users would strategically delay this upgrade.

The third research question (RQ3) seeks to causally identify the outcomes associated with delaying upgrade. Specifically, I ask and answer, *"How does delaying upgrading affect apps' marketplace outcomes (app rating and app popularity)".* Here,

56

I focus on the effect of timing of mobile apps' decision to upgrade to Android 6.0. From the literature survey, it is unclear whether being an early-mover in a policy adoption scenario is advantageous at all. Understanding if there are repercussions to strategically delaying upgrading, has important policy implications for an open platform such as Android. The next section describes my research context and data to address the above-mentioned questions.

3.3   Research Context and Data

I examine my research questions on the app marketplace (called Play Store) of Android, the world's largest smartphone platform (Statista 2018a). Despite Android launching version 6.0 in late 2015, my panel starts from April 2016 when the user installation base for Android 6.0 reached 5% (Android 2016, Statista 2018b). The choice of time-line addresses the potential lack of motivation among apps in the initial months when there aren't many Android users who have upgraded to Android 6.0[22]. I compile my data by combining two data sources that I gather over two years.

First, I take a list of 21,000 [23] most downloaded free/freemium apps as of April 2016 and install these apps (over 1 terabyte) on emulators on research PCs at my lab. Installing apps is the only clean way to extract the Android version number of that app (*targetAPI*) since only an installed file carries the signature (SDK version) with the *targetAPI*. I chose free apps instead of paid apps because paid apps tend to be less strategic in terms of permissions, due to lower reliance on user's information for

---

[22] Initial 5% of users have been shown to exhibit willingness to take risks with new features and provide valuable feedback to products for further improvement (Mahajan and Muller 1998), thereby helping app evolution.

[23] The list contains all apps that had greater than 500,000 downloads (Around 18,600 apps) and about 2400 apps with 100,000-500,000 or lower downloads.

revenues. I developed and installed a custom-built Android app that scans installed apps and extracts the *targetAPI* of each installed app, app version number, and the list of all permissions sought by those apps. I updated all installed apps monthly to collect the data points mentioned above for 24 months.

My second data source is the Android's Play Store website from which I exhaustively collect details of over two million Android apps on a monthly basis. I collect information such as app description, rating, download count bucket, count of reviews, categories, date of last update, number of screenshots uploaded, app's file size (in MB), developer ID of all apps and two dummy variables, one each to indicate whether the app has *in-app advertising* or *in-app purchase* options in for each month. Merging this dataset with the earlier dataset gives me an unbalanced panel of 13,604 [24] apps for 24 months (April 2016 to March 2018), resulting in 278,955 app-month observations[25].

3.3.1 Essential and Non-essential Permissions

A key task in my investigation is to determine which of the permissions sought are essential for the app's working. To achieve this, I divide apps into sub-categories (called *peer groups*) of similar functionality based Peer Group Analysis Technique (Pelikan et al. 2017). I then use a statistical technique similar to Permission Prevalence Analysis (Taylor and Martinovic 2016) to determine which permissions are *non-essential* for app's working. First, I employ Skip-gram Word2Vec, a modern and highly

---

[24] Of the list, 6648 apps were not available to download on emulators in the United States. Another 748 apps got pulled out of the Play Store with less than 12 months in my panel resulting in 13,604 unique apps in my panel.

[25] While we should have 326,496 observations (13,604*24) in my dataset, some of the apps are not available in Play Store for all 24 months, mostly because they are pulled out of US Play Store.

effective text-mining technique, to obtain vector representation for each app. A distributed vector representation of words, that learns the locational similarity and context of words in a statement, has been proven more accurate than the traditional bag of words or n-gram vectors (Mnih and Hinton 2009). For my analysis, this technique is preferred over Term Frequency Inverse Document Frequency (TF-IDF), because of Word2Vec's ability to capture synonyms, contexts of words or grammatical nuances such as plurals (Ramos 2003). Furthermore, TF-IDF may suffer from *curse of dimensionality* when the corpus is large (Hinneburg and Keim 1999). Word2Vec by Mikolov et al. (2013) resolves these issues by using Skip-gram model to create a 300-dimensional distributed vector representation for each word. I get a single 300-dimensional vector for a document by combining individual word vectors of that document. Since app descriptions are written in standard English with carefully curated English words/phrases, I utilize a model trained on 100 billion words from a Google news dataset to improve the accuracy of dividing apps into peer-groups based on functional similarity.

I start with tagging each app description from Play Store with a category label determined by Android (see Appendix Table A6 for a comprehensive list of categories). Since these categories are pre-determined by Android, I am assured of the general similarity between apps under the same category. Next, I use an unsupervised learning technique to divide apps within each category into sub-categories of apps that have similar functionality and utility.

Codifying the process to determine the sub-category of apps, I:

I. Obtain a 300-dimensional vector for each app based on app description by employing Word2Vec.

II. employ k-means clustering to determine the optimal sub-categories within each category
   a. Determine the range of optimal clusters [k_m-k_n] for app category c using Elbow technique, AIC, BIC
   b. Determine the least number of clusters (k_c), within [k_m-k_n], wherein each cluster has at least 5 apps.
   c. Divide apps under each category c into k_c clusters (i.e., k_c sub-categories of each of c category)

To illustrate the effectiveness of sub-categorization, I chart a word-cloud for the sub-categories created under two of the most generic categories, "Productivity" and "Lifestyle". As seen in Figure 3.2, the sub-categories are quite homogenous while distinct from each other. I also manually verify a random sub-set of all algorithmically generated sub-categories to ensure that the sub-categorization has worked well.

Next, similar to the statistical technique proposed by Sarma et al. (2012) or Taylor and Martinovic (2016), I determine which of the permissions sought are essential to all apps in each sub-category/peer-group. I code those permissions requested by more than 75% of the apps in a sub-category/peer group as *essential* permissions and the rest as *non-essential* permissions. To build intuition, I consider navigation apps: permission to access GPS is essential for any navigation app; hence, at least 75% of such apps would seek permissions to GPS. On the contrary, access to users' call log does not seem essential to navigation apps; hence, I would expect less than 75% of the apps to seek permission to access the call log. I found qualitatively similar results by varying the threshold between 60% and 90%.

3.3.2 Popularity of Apps

Play Store does not publicly disclose the exact download count for each app. This precludes researchers from using download counts to estimate an app's popularity. Hence, I determine app popularity by building a network of apps and estimating each

app's network centrality. For this, I employ Weighted PageRank algorithm (Page et al 1999) to derive app popularity score. Researchers (Oestreicher-Singer and Sundararajan 2012; Kane and Ransbotham 2016) have employed variants of PageRank algorithm to derive the importance of various nodes in a network using the concept of *eigenvector centrality*. Intuitively, an influential node has a larger *eigenvector centrality*. Alternatively, a larger value of *eigenvector centrality* for a node also means that the node is pointed to by (i.e., receives a vote of support from) many nodes. I employ the latter interpretation of *eigenvector centrality* in my investigation. For each app in the Play Store, Android provides a list of "similar" apps (see Figure 3.3).

The similar apps list can be thought of as a "vote of support" from the focal app to the similar apps. Intuitively, the greater the popularity of an app, the greater the likelihood that such an app would appear on similar apps' list of many apps in a given category. Therefore, by gathering similar apps' lists for each of the apps, I can estimate how popular the app is. I employ a modified variant of PageRank algorithm, where I use an app's position in the similar apps list to determine the strength of the vote of support. Specifically, I use the formula *max(1-rank/24, 0.25)* to determine the strength of the directed edge from the focal app to similar apps. I repeat this process for each month to estimate app's popularity in any given month.

3.3.3 Variable Definitions

Table 3.1 provides a detailed description of the variables used in the major analyses. Appendix Tables A7 and A8 provide the summary statistics and correlation matrix of the variables.

Variables of Interest: Our first three variables of interest are defined as the count of total permissions sought, essential permissions sought, and non-essential permissions

sought, respectively. We normalize the three variables by computing the ratio of permissions sought to all permissions utilized by the app (dangerous and normal[26]) and use the three ratios as independent variables in survival models.[27] Variable $upgrade_{it}$ is coded as 1 once the given app performs an upgrade to the latest version and 0 before that. Upon upgrading from version 5.1 (or below) to 6.0 (or above), the value changes to and stays at 1. Finally, $rating_{it}$ captures the user rating that the app i has received by the end of month t. Finally, we derive $popularity_{it}$ using the weighted PageRank algorithm as outlined in Section 3.2.

Explanatory Variables: The explanatory variables are app characteristics that represent app demand, apps' appearance on the Play Store and developers' details from Play Store. For each app on a given month, I capture the total count of ratings provided by users, number of days elapsed since the developer has pushed an update (features, bug fixes or upgrades), total count of apps by focal app's developer on Play Store, count of screenshots that the app has on app's Play Store, file-size and revenue model (*in-app-purchases* or *in-app-advertising*).

Control Variables: I use a set of control variables to account for unobserved app characteristics. I classify the app categories into six groups based on prior research (for example, (Ghose and Han 2014)) and how categories evolve over time. The six categories include: Online Content Consumption (Media and Entertainment), Learn and Explore, Personal (Social and Lifestyle), Mobile Specific Utilities (services that

---

[26] Normal Permissions cover areas which pose very little risk to user's privacy (e.g., setting alarm, accessing Wi-Fi state). Consumers don't have a choice to rejecting them since developers are not required to explicitly seek permissions from users.
https://developer.android.com/guide/topics/permissions/overview.html#normal-dangerous
[27] We re-estimate all models with count variables (instead of ratios) and find qualitatively consistent results.

are enabled by a cellphone technology such as GPS), Mobile Access Utilities (mobile apps provided by offline or internet utility firms), and Games. Finally, $download\_bucket_{it}$ captures the range of app downloads that Play Store publishes[28].

3.4.   <u>Analyses and Results</u>

3.4.1   What are the characteristics of apps that delay upgrading?

To investigate how an app's characteristics and its permission-seeking behavior affects the app's decision to delay upgrade to the latest version, I employ a continuous-time single-failure survival model. Specifically, I employ a Cox Proportional-Hazard model that investigates how various covariates of interest affect upgrading to the latest version of Android. In my analysis, I consider upgrading to Android 6.0, represented by a dichotomous variable, as *failure*. The survival model is specified as below:

$$\lambda_i(t) = \lambda_0(t) \exp(c_1* sought\_permissions\_ratio_i + c_2*\Delta rating\_count_i + c_3* rating_i + c_4* dayssinceupdate_i + c_5*screenshots_i + c_6* filesize_i + c_7*developer\_appcount_i + \Gamma*download\_bucket_i + \Lambda*category\_group_i) \quad (3.1)$$

where the subscript i indexes apps and subscript t indexes months, $\lambda_i(t)$ is the hazard of the app i upgrading to Android 6.0 at time t while $\lambda_0(t)$ represents the baseline hazard of upgrading to Android 6.0. Single-failure specification discards the observations after *failure* has occurred. For each of the covariates, a hazard rate less than 1 (HR<1) indicates that the co-variate decelerates the chances of failure (in this case, upgrading to Android 6.0), i.e., increases the chances of delaying upgrade and vice versa with HR>1. Such an interpretation is consistent with literature (for example, (Kauffman et al. 2000)). In this model, I pool *essential* and *non-essential* permissions in order to

---

[28] The download buckets: 1-5 downloads, 5-10, 10-50, 50-100, 100-500, 500-1k, 1k-5k, 5k-10k, 10k-50k, 50k-100k, 100k-500k, 500k-1M, 1M-5M, 5M-10M, 10M-50M, 50M-100M, 100M-500M, 500M-1B and 1B-5B downloads

focus on the effects of app characteristics and apps' marketplace standings on delaying upgrade. I explore the heterogeneous effects on later models. Outcomes of model (3.1) are presented in Table 3.2.

Column 1 provides the point estimates while column 2 provides the corresponding Hazard Ratio (HR). The hazard rate of $rating_{it}$ is significant and greater than 1, indicating that apps with higher ratings are less likely to delay upgrade. Apps that have higher monthly review counts, an indicator of app's efforts to solicit ratings and feedback, have a lower propensity to delay upgrade. The number of days since the latest update, an indicator of app's maintenance cycle, increases the propensity to delay upgrade. App file-size, an indicator of the sophisticated nature of the app and number of screenshots uploaded to the Play Store, an indicator of app demand (Ghose and Han 2014), negatively impact the propensity to delay upgrade. Finally, I find that apps that have less than 1-million downloads are significantly more likely to delay upgrade to the newer Android version while apps with more than 10 million downloads are less likely to delay upgrade to the latest version.

Shifting the focus to app characteristics, I find that, compared to the "online content consumption" apps (such as video streaming services), mobile-specific utility apps (such as the messaging, navigation, food delivery apps) and game apps are more likely to delay upgrade to the new version of Android. Similarly, personal apps (such as social networking/dating apps) are more likely ($p<0.10$) to delay upgrade. This is consistent with reports that have found that apps from the above-mentioned categories are more likely to over-seek *non-essential* permissions (Maheshwari 2017, Stamm 2018).

64

I now focus on effects of permission-seeking behavior on upgrade decisions. I find that an app's likelihood of delaying upgrade to the latest Android version increases with an increase in the ratio of permissions sought by the apps. This indicates that apps that seek more permissions prefer to retain control over access to users' private information for as long as they can.

To further investigate how seeking individual permissions affect developers' propensity to delay upgrade, I alter the Cox PH model in (3.1) to the following:

$$\lambda_i(t) = \lambda_0(t) \exp(C * permissions\_vector_i + c_2 * \Delta rating\_count_i + c_3 * rating_i + c_4 * dayssinceupdate_i + c_5 * screenshots_i + c_6 * filesize_i + c_7 * developer\_appcount_i + \Gamma * download\_bucket_i + \Lambda * category\_group_i) \tag{3.2}$$

where the subscript i indexes apps and t indexes months. Here, $permissions\_vector_{it}$ is a vector of dummies, indicating the presence of individual permissions and C is a vector of coefficients for this vector of dummies. Table 3.3 carries the results of the estimation of model (3.2). As seen in Table 3.3, only specific permissions that users may consider as sneaky or running in the background, such as permissions to access users' information (call logs, phone memory, accessing user's ID or user's phone status) or their phone hardware (microphone or fine GPS location) increase the propensity to delay upgrade to Android 6.0 (i.e., HR<1 and significant).

In summary, the analysis indicates that seeking more permissions, increases the propensity to delay upgrade. In the next section, I perform multitudes of tests including a falsification test and forward arguments to rule out alternative explanations.

3.4.1.1 Ruling out Alternative Explanations

An app's decision to delay upgrade may be due to the cost related to upgrading, the lack of demand for an upgrade from the existing users, the lack of benefit from

upgrading or the app's desire to retain control over users' information. In this section, I rule out these alternative explanations.

Developers incur a cost to upgrade because the process of upgrading to a latest version requires them to understand changes in the SDK (such as new features) and modify app control flow (if needed) based on best practices recommended by the latest SDK. Such cost incurred may heterogeneously affect developers depending on their size/capacity. I account for such heterogeneity in my models by controlling for developers' monthly app count (indicator of their size) on Play Store. Second, by focusing my analyses on the most-downloaded apps, I further mitigate cost-side concerns. This is because I expect developers of popular apps to have the motivation and monetary incentives to keep their apps updated and relevant. Next, I account for apps' sophisticated nature and maintenance cycle by including the app characteristics such as its file size, popularity (i.e., count of reviews left by app users) and maintenance cycle ($dayssinceupdate_{it}$) in the models. Finally, in unreported results, I estimate the Cox Proportional Hazard model with developer fixed effects to investigate if the results hold. Using developer fixed effects discards observations of 5,073 apps where none of the apps by a developer upgrades to Android 6.0 during the duration of the panel. The results are consistent with the main analysis, further confirming that my results are not driven by developer heterogeneity. Another major reason why apps might delay upgrading may stem from user demand. Apps may have lesser incentive to upgrade if they cater to a majority of users that will not upgrade their phones to the latest Android version. While such a scenario (existing users not upgrading) is feasible, Android's platform strategy provides good reasons why those apps should still upgrade. Android

66

maintains a "Quality Score" for listed apps in Play Store[29] and explicitly states that the quality score, and therefore the app's Play Store ranking, may be affected by delaying upgrade. A better Play Store rank aids in better discovery in the Play Store when new users search for a specific utility or when they scroll through the "similar apps" list. Furthermore, despite upgrading to the latest Android version, apps would still seek download-time permissions to user information from users that do not upgrade. Hence, platform's design decisions help rule out demand-related reasons for delaying upgrade.

Another explanation for why certain apps may have no incentives to upgrade is the lack of benefits (to apps and/or users) from upgrading to the latest android version. Simply put, only apps that find new Android features more beneficial are more likely to upgrade while others do not invest effort to upgrade. A survey of all changes incorporated in the 6.0 version of Android informs me that all new features and improvements would be beneficial to all apps equally. For example, Android made significant improvements to its notification services, which would equally improve the user notification effectivity for all apps. Furthermore, I also find that none of the changes made to existing features in Android (other than changes to permission seeking) adversely affects apps' functionality. Hence, I am confident that there are no other feature improvement motivations that may impact apps' strategies for delaying this upgrade. Thus, choosing not to upgrade based on lack of benefits from upgrading does not appear to be the driving reason.

3.4.1.2 Falsification Test: Upgrading to an earlier version of Android

---

[29] https://developer.android.com/docs/quality-guidelines/core-app-quality

In the earlier section, I argue that the ability to continue seeking (potentially unnecessary) permissions is the primary reason why apps delay upgrading to Android 6.0 (*targetAPI* 23). A falsification test is to check whether such a finding holds/fails when apps upgrade to an earlier version (i.e., prior to *targetAPI* 23). If losing the ability to seek blanket permissions were indeed the primary reason for delaying upgrade to *targetapi* 23, then seeking more permissions should not increase the propensity to delay upgrading to earlier versions, such as Android 5 (or Android 5.1). To test this, I employ a Cox Proportional Hazard model similar to model (3.1) with a few changes. In this model, *failure* is determined by a dichotomous variable that carries the value 1 when and after the app upgrades to *targetAPI* 22 (Android 5.1) or above from a lower version. Consistent with the earlier approach, observations post-failure are discarded in the Cox Proportional Hazard model. Table 3.4 carries the results of the falsification test.

In both columns (1) and (2) of Table 3.4, the Hazard Ratio of *sought_permissions_ratio$_{it}$* is significant and greater than 1. This differs my findings from Table 3.2, where I found the hazard ratio to be significant and less than 1. Upon carefully investigating new features that were introduced in *targetAPI* 21 or *targetAPI* 22, I find that in Android 5.0 and 5.1, the Android platform revoked restrictions (which Android had enforced in an earlier version) that prevented apps from accessing files that did not belong to the app's own home directory[30]. Android also introduced screen capture and sharing functions, as well as features to programmatically access users' camera devices (new camera API), app usage history, and battery usage logs. These

---

[30] https://developer.android.com/about/versions/android-5.0

features incentivized apps seeking more (non-essential) permissions to upgrade to *targetAPI* 21/22 but not *targetAPI* 23.

3.4.2 Is the app's decision to delay upgrading to seeking run-time permissions strategic?

This research question builds on RQ1 by investigating whether apps' decision to delay upgrade is linked to apps' monetization from continued access to non-essential (to app's core functionality) sensitive information. To address this research question, I replace the variable *sought_permissions_ratio$_{it}$* in the model (1) with variables that measure the *essential* and *non-essential* permissions respectively. As outlined in section 3.1, I use Skip-gram Word2Vec and k-means clustering to sub-categorize Android apps into *functionally similar* groups and then statistically distinguish contextually *essential* permissions from *non-essential* permissions.

$$\lambda_i(t) = \lambda_0(t) \exp(c_1 * essential\_permissions\_ratio_i + c_2 * \Delta rating\_count_i + c_3 * rating_i + c_4 * dayssinceupdate_i + c_5 * screenshots_i + c_6 * filesize_i + c_7 * developer\_appcount_i + \Gamma * download\_bucket_i + \Lambda * category\_group_i)$$ (3.3)

$$\lambda_i(t) = \lambda_0(t) \exp(c_1 * nonessential\_permissions\_ratio_i + c_2 * \Delta rating\_count_i + c_3 * rating_i + c_4 * dayssinceupdate_i + c_5 * screenshots_i + c_6 * filesize_i + c_7 * developer\_appcount_i + \Gamma * download\_bucket_i + \Lambda * category\_group_i)$$ (3.4)

where the subscript *i* indexes apps and *t* indexes months. If both types of permissions types (i.e. *essential* and *non-essential* permissions) equally affect apps' decision to delay upgrade, I expect to see similar Hazard Ratios for the coefficients in models (3.3) and (3.4) as in model (3.1) (i.e., HR<1 and significant for c_1). A change in HR (from HR<1 to HR>1) or different statistical significance of one of these two coefficients would mean that there is a heterogeneous effect of *essential* and *non-essential* permissions on apps' propensity to delay upgrade. The results are shown in Table 3.5.

The hazard ratio in column (1) for *essential_ permissions_ratio$_{it}$* is significant and greater than 1 while the hazard ratio in column (2) for *nonessential_ permissions_ratio$_{it.}$* is significant and less than 1. Taken together, these indicate that the propensity to delay upgrading to the latest version is influenced by only the *non-essential* permissions sought by apps. This finding is consistent with my earlier argument that apps that over-seek *non-essential* permissions are more likely to delay upgrade.

3.4.2.1 What drives apps to over-seek *non-essential* Permissions?

My investigation so far suggests that apps that overseek *non-essential* permissions delay upgrading to Android 6.0. A natural question that arises from my investigation is, *why do some apps seek non-essential permissions, which in-turn increases their propensity to delay upgrading?* The results in Table 3.3 shows that permissions to access users' personal information such as call logs, phone memory, user ID or phone status or permissions to access their hardware such as microphone or GPS increases the propensity of apps to delay upgrade. Apart from the sneakiness of these permissions, reports have shown that users' sensitive information, such as GPS location, microphone recordings or personal connections, have been abused by apps for personalized ad-targeting (Goode 2018, Limer 2018). It stands to reason then, that apps that rely on users' sensitive information to generate revenue are more likely to over-seek *non-essential* permissions. Such over-seeking behavior should in-turn drive these apps to strategically delay upgrade so that they continue collecting sensitive information for as long they can. In unreported analysis, I utilized machine-learning techniques to further support my intuition that the decision to delay upgrade are driven

70

by apps' ad-targeting behavior. Specifically, I employ a combination of SkipGram based Topic Modeling and Lasso Logit to identify topics that predict delaying upgrade. I find high similarities (cosine similarity) between these topics and media report clippings that report on aggressive advertising or "adware" in Android.

I now investigate whether the app's revenue model impacts its likelihood of delaying upgrading to Android 6.0. Of the two main freemium revenue models in Play Store, *in-app-advertising (IAA)* and *in-app-purchases (IAP),* I expect apps with an *IAA* revenue model to have stronger incentives to seek more permissions and delay upgrade. Alternatively, I expect apps with *IAP* revenue model not to impact the decision to delay upgrade. As expected, Appendix Table A9 shows that *IAA* revenue model positively impacts the link between permissions and decision to delay upgrade while *IAP* revenue model does not (neither directly nor indirectly) impact the decision to delay upgrade.

To rigorously test the link between *IAA* revenue model and the propensity to delay upgrade, I follow a two-stage control function approach called the Two-Stage Residual Inclusion (2SRI) method (Terza et al. 2008, Wooldridge 2015). The procedure involves a two-stage model wherein the first stage estimates the effect of *IAA* revenue model on seeking permissions while the second stage estimates the effect of seeking permissions on the propensity to delay upgrade. Since the second stage is a survival model, following Terza et al. (2008), and Wooldridge (2015), I use the 2SRI approach instead of the traditional 2SLS approach. Such a procedure has been widely used across social science disciplines including Finance (Chen et al. 2013), Marketing (Risselada et al. 2014), Sociology (Lyons et al. 2013) and Economics (Beaudry and Allaoui 2012). I

71

employ a generalized version of 2SRI model that uses individual frailty (Martínez-Camblor et al. 2017, Martínez-Camblor et al. 2019).

The decision of utilizing *IAA* revenue model should affect apps' decision to over-seek *non-essential* permissions for sensitive information. However, unlike other revenue models, including display advertisements within apps neither increases the complexity of app nor requires extensive feature development. Therefore, displaying *in-app-advertising* should not operationally affect apps' decision to delay upgrade, except through their decision to over-seek *non-essential* permissions. Hence, the excluded instrument from the first stage satisfies the exclusion restriction. The statistical tests (i.e., under-identification test and weak first stage test) suggest that the excluded instrument is statistically valid. Table 3.6 shows the outcome of this estimation.

From column (2) of Panel (A) find that having *in-app-advertising (IAA)* as a revenue model has a strong and significant effect on seeking non-essential permissions, while IAA does not impact seeking *essential* permissions (see Column (1) of Panel A). the decision to delay upgrade, but only via its impact on *non-essential* permissions. The statistical tests (i.e., under-identification test and weak first stage test) suggest that the excluded instrument is statistically valid for *non-essential* permissions model.

The second stage of both the models (Panel B of Table 3.6) show qualitatively consistent results as Table 3.5, ascertaining my claim that apps that overseek *non-essential* permissions delay upgrading.

To visualize the findings, I plot the *essential* and *non-essential* permissions sought by apps across the Android versions over three different period spread across the panel.

As seen in Figure 3.4, there is a significant drop in seeking *non-essential* permissions moving from *targetAPI* 22 to *targetAPI* 23, but insignificant changes in *essential* permissions, further adding strength to my arguments.

3.4.3 Effect of timing of upgrading on the outcome for apps

I seek to understand if Android's privacy policy change is indeed effective in improving apps' sensitive information-gathering behavior by investigating the policy's effect on apps' permission-seeking behavior. It is also useful to learn whether strategically delaying adhering to such policy changes are accompanied by marketplace penalties, as measured by the app ratings AND app's popularity on Play Store. To address the questions on the impact of apps' time choice of upgrading, I employ Propensity Score Matching (PSM) to pre-process the data (Dehejia and Wahba 2002) and employ a Difference-in-Differences technique (DiD) to establish causality of the relationship between the treatment and the treated (Meyer 1995).

3.4.3.1 Model Setup

The timing of upgrading to the latest version is a choice that apps exercise; hence the process of upgrading is considered a *treatment* in my setup. A DiD approach enables me to compare and contrast the effects on apps that upgrade to the latest versions (the *treated* apps), with those apps that do not upgrade until the end of my panel (the *control* apps). Since the *treated* apps upgrade over a period of several months, i.e., in a staggered manner, I set the month of treatment (month when the *treated* app upgrades) to time 0 and adjust the months prior to and post that in reverse chronological order (…-3, -2, -1) and sequential order (1, 2, 3…) respectively. Such a normalization of time is performed based on prior research (Autor 2003, Fang et al. 2014, Gao and Zhang 2016). Employing DiD estimation approach in a non-experimental setup such as mine

requires that I correct for treatment selection bias. To address this, I employ PSM as my pre-processing technique. PSM has been used in many non-experimental settings (Dehejia and Wahba 2002). Specifically, techniques combining PSM to pre-process the dataset followed by DiD have been used in research where the treatment may be affected by a selection bias (Liu and Lynch 2011, Smith and Todd 2005). If matched correctly, the *treated* and *control* samples will have the same propensity to upgrade to the latest version in a given month, with the decision to upgrade being the only difference. In other words, each *treated* app (an app that upgrades within the time-period of the panel) is matched with a *control* app that never upgrades until the end of the panel. I apply nearest one-to-one neighbor matching without replacement of the *control* samples. To account for the difference in time periods when treatments are introduced (app A upgrades in June 2016, app B upgrades in July 2016), I follow a dynamic matching technique where I match the *treated* apps with *control* apps one month before the month of treatment.

The covariates used for matching include app characteristics such as the app rating, count of rating, file size, screenshots uploaded, date since the latest update and revenue model dummies. Also, I explicitly ensure that *treated* and *control* apps come from the same download bucket and belong to the same app category. The matching procedure gives me 2628 *treated* and 2628 *control* mobile apps in the final sample. As the first test for successful matching procedure, I compare continuous covariates in my models before and after the matching procedure. The outcomes of covariate balance tests for the PSM technique is presented in Table 3.7. The covariates are well balanced after matching procedures.

Model (3.5) estimates whether, and by how much, upgrading to the latest version changes apps' information gathering behavior, reflected by permissions they continue to seek. Model (3.6) estimates the impact of upgrade on their apps' marketplace outcomes as seen from apps' ratings.

$$\textbf{\textit{essential\_permissions}}_{it} = \alpha_{0a} + \alpha_{1a}*\Delta rating\_count_{it} + \alpha_{2a}*rating_{it} + \alpha_{3a}*post\_upgrade_t + \alpha_{4a} *upgrade\_group_i*post\_upgrade_t + \alpha_{5a}* dayssinceupdate_{it} + \alpha_{6a}*filesize_{it} + \alpha_{7a}*screenshots_{it} + \alpha_{8a}* developer\_appcount_{it} + \alpha_{9a}*download\_bucket_i + \alpha_{10a}*month\_dummy_t +w_i + \varepsilon_{ait}$$

(3.5a)

$$\textbf{\textit{nonessential\_permissions}}_{it} = \alpha_{0b} + \alpha_{1b}*\Delta rating\_count_{it} + \alpha_{2b}*rating_{it} + \alpha_{3b}*post\_upgrade_t + \alpha_{4b}*upgrade\_group_i*post\_upgrade_t + \alpha_{5b}* dayssinceupdate_{it} + \alpha_{6b}*filesize_{it} + \alpha_{7b}*screenshots_{it} + \alpha_{8b}* developer\_appcount_{it} + \alpha_{9b}*download\_bucket_i + \alpha_{10b}*month\_dummy_t +w_i + \varepsilon_{bit}$$

(3.5b)

$$\textbf{\textit{popularity}}_{it}= \beta_{0a} + \beta_{1a}*\Delta rating\_count_{it} + \beta_{2a}* sought\_permissions\_ratio_{it-1} + \beta_{3a}* rating_{it} + \beta_{4a} * post\_upgrade_t + \beta_{5a}*upgrade\_group_i*post\_upgrade_t + \beta_{6a}* dayssinceupdate_{it} + \beta_{7a}*filesize_{it} + \beta_{8a}* screenshots_{it} + \beta_{9a}* developer\_appcount_{it} + \beta_{10a}*download\_bucket_i + \beta_{11a}*month\_dummy_t +w_i + \eta_{ait}$$

(3.6a)

$$\textbf{\textit{log(rating}}_{it})= \beta_0 + \beta_1*\Delta rating\_count_{it} + \beta_2* sought\_permissions\_ratio_{it-1} + \beta_3* post\_upgrade_t + \beta_4*upgrade\_group_i*post\_upgrade_t + \beta_5* dayssinceupdate_{it} + \beta_6*filesize_{it} + \beta_7* screenshots_{it} + \beta_8* developer\_appcount_{it} + \beta_9*download\_bucket_i + \beta_{10}*month\_dummy_t +w_i + \eta_{it}$$

(3.6b)

In the above models, $upgrade\_group_i$ is a binary variable that carries a value of 1 for apps in the *treated* group and 0 for apps in the *control* group. For each treated-control pair, the variable $post\_upgrade_t$ carries a value of 0 before the *treated* app performs an upgrade and 1 after it upgrades to the latest version. Since I follow the time normalization procedure described above, $upgrade\_group_i*post\_upgrade_t$ identifies the effect of upgrading as this coefficient equals 1 only for treated apps ($upgrade\_group_i$ =1) post the upgrade ($post\_upgrade_t$=1). Coefficients α_{4a}, α_{4b} and β_4 in the above models above provide me with required DiD estimates for the model. The variable w_i represents the app-level fixed effects to account for any app level

unobserved heterogeneity. Logging the rating variable and scaling the variable by 100 allows me to interpret the outcome as percentage change. While I separately estimate the outcomes on *essential* and *non-essential* permissions in the main analyses, I conduct robustness checks by pooling the permissions.

Parallel trends assumption states that the control and treated apps should have a similar time trend before the treatment. To test this assumption, I follow the recommendations in literature (Autor 2003) where I estimate models similar to (3.5) and (3.6), but by modeling time trends of the dependent variables in the models. This can be done by introducing interaction terms of the treatment and the pre-treatment month dummies. The models are quite similar to equations (2.5) through (2.7) in the previous chapter.

I list the outcome of these models in Appendix Table A10. All interaction terms before the month of upgrade are insignificant, therein assuring that the Parallel Trends Assumption is not violated by the matching technique. Figure 3.4 provides the pictorial representation of parallel trends assumption for essential and non-essential permissions.

3.4.3.2 Determining if apps are delaying upgrading to Android 6.0

Apps indulging in strategic delaying behavior is neither beneficial to the customers (they lose out on enhanced security, improved reliability, better UI and UX) nor to the platform (because it would exacerbate the existing issue of supporting multiple SDK versions). Hence, I investigate if strategic delaying has repercussions on apps' outcomes. To study the effect of timing on apps' information gathering behavior and marketplace outcomes, I distinguish early/on-time upgraders from late upgraders

among apps in my models. Since Android launches a new version every year, I consider all those apps that upgraded to Version 6.0 (or above) after users started to upgrade to Version 7.0 of Android as late upgraders among apps. While Android announced their new version (Version 7.0) at the end of August 2016, the Android 7.0 customer base crossed 0.1% of overall Android customer base only by November 2016. Hence, I consider all apps that upgraded from version 5.1 or below (download-time permissions) to 6.0 or a higher version (run-time permissions) in and after November 2016 as late upgraders. I alter models (3.4) and (3.5) to incorporate another two-way interaction term between the dummy *late_upgrade_group$_i$* that carries the value of 1 if the *treated* app upgrades after November 2016 and the dummy *post_upgrade$_t$* that indicates whether the current time-period is before or after the treatment in DiD models.

3.4.3.3 How does delaying upgrading to seeking run-time permissions affect the outcome for apps?

In this section, I investigate whether apps' decision to upgrade are accompanied by their change in information gathering behavior as well as marketplace outcomes.

3.4.3.3.1 Does delaying upgrading affect apps' sensitive information gathering behavior?

Model (3.4) estimates the effects of upgrading on app's information seeking behavior, the outcomes of which have been presented in Panels A and B of Table 3.8. Columns (1) and (2) present the DiD estimation of upgrading while columns (3) and (4) present the estimation by including an interaction term between *late_upgrade_group$_i$* and *post_upgrade$_i$*. This interaction term identifies the effect of delaying the upgrade on the dependent variables, over and above the effect of upgrading to the latest version of the

app (which is captured by the interaction term between $upgrade\_group_i$ and $post\_upgrade_t$).

From column (3) of Panel A, I find that delaying upgrading to a new platform version does not impact the *essential* permissions sought by apps, as seen by coefficient of $late\_upgrade\_group_i*post\_upgrade_t$, which is insignificant. On the other hand, from column (3) of Panel B, I find that apps significantly reduce seeking *non-essential* permissions. Delaying upgrade reduces apps' ability of over-seeking *non-essential* permissions by 6.79% (i.e., {0.101 + 0.083}/2.71). This means that the apps' ability to over-seek *non-essential* permissions further falls upon delaying upgrade. Taken together, these findings suggest that apps significantly reduce their information over-seeking behavior upon delaying upgrade.

3.4.3.3.2 Does delaying upgrading affect apps' marketplace outcomes?

Model (3.5) estimates the impact of upgrading on marketplace outcomes, the outcomes of which have been presented in panels C and D of Table 3.8. Upgrading to Android 6.0 on time positively affects apps' marketplace outcomes, as seen by the positive and significant result of coefficient $upgrade\_group_i*post\_upgrade_t$ in column (3) of Panel C. This coefficient (0. 16822) means that, on an average, about 1020 app users (i.e., 0. 16822/100* 4.04 mean rating * 150079.3 mean raters) who have previously rated the app, increase their rating by 1 star after the app upgraded to the latest version of the platform. However, delaying upgrading has a cost to pay for the apps. The coefficient $late\_upgrade\_group_i*post\_upgrade_t$ in column (3) suggests that delaying upgrading erases the potential rating gain that apps would have received. The coefficient of - 0.09237 means that delaying upgrading results in a fall in rating, an equivalent of 560

users reducing the rating by 1 star (i.e., 0. 09237/100* 4.04 mean rating * 150079.3 mean raters).

While delaying upgrade adversely affects apps' ratings, it is unclear whether this adverse effect of delaying upgrade on rating are towards all apps or only those apps that continue to over-seek *non-essential* permissions. To further understand the reasons for this adverse effect, I perform sub-sample analyses based on the changes in information seeking behavior of *treated* apps. I divide treated apps into two groups based on whether apps' permissions count falls more (or less) on average after upgrade as compared to before upgrade. Appendix Table A11 shows the outcomes of the sub-sample analysis. From column (1) I find that the adverse effect of delaying upgrade on ratings are only towards apps that do not reduce over-seeking *non-essential* permissions after upgrade. This further supports my conjecture that strategically delaying upgrade, because of permission over-seeking behavior, also has rating repercussions.

Finally, Panel D of Table 3.8 provides the outcomes of Model (3.6B), i.e., the impact of delaying upgrade on apps' popularity. From column (3) of Panel D, I find that delaying upgrade significantly negatively affects the popularity by 9.06% (0.822/9.07, mean value of the popularity). I find consistent results with column (4) where I present weighted regression. In essence, strategically delaying upgrading to the latest version of Android has both monetary and reputational repercussions, while forcing the apps to significantly improve their sensitive information gathering behavior.

3.4.3.4 Additional Robustness Checks

Since my data-collection process involves installing apps on emulators, there is potential oversampling of apps based on its popularity. Despite using a large sample of apps and employing global app distribution as probability weights in regression analysis, questions regarding generalizability of my finding may still exist. As an alternative test, I divide the apps into 3 groups based on the overall download categories: less than 1 million downloads (Low Download Group, $lowDG_i$), 1 million – 5 million downloads ($mediumDG_i$) and greater than 5 million downloads ($highDG_i$)[31]. I re-estimate models (5) and (6) by including a 3-way interaction term *late_upgrade_group$_i$*post_upgrade$_t$*download_group$_i$*, wherein the baseline comparison group is the *highDG*. The results are shown in Appendix Table A12. I see that the outcomes of delaying upgrading to Android 6.0 on apps' overseeking behavior (i.e., overseeking *non-essential* permissions) as well as on apps' ratings are most conservative for apps in *highDG*, followed by apps in *mediumDG* and subsequently for those in *lowDG*. This indicates that the negative impact of delaying upgrading on apps' rating as well as apps' permission seeking behavior only exacerbates as I include apps with lower downloads.

3.5.   Discussions

3.5.1   Key Findings

This study is motivated by an increasing demand by users to take control over their sensitive information and a dearth of research examining the impact of platforms' mechanisms that seek to empower users. I investigate app's responses to Android's

---

[31] An app might move within groups over 24 months. We categorize such apps into the group where they spent most of their time in. E.g., spending 8 months in *lowDG* and 16 months in *mediumDG* puts an app in *mediumDG* category

upgrade that changed how Android users grant permissions to apps. In investigating the characteristics of apps that delay upgrade, I find that the propensity to delay upgrade falls for high quality apps, apps that are characterized by high consumer satisfaction (based on user ratings), a really active consumer base (based on number of comments/feedback posted each month), an active maintenance cycle (based on the days since last update), sophisticated features (based on larger file size), high app demand (based on screenshots) or large download base. These characteristics are consistent with literature on high quality software/mobile apps (Ghose and Han 2014, Jung et al. 2012, Krishnan et al. 2000).

A key finding of my investigation is that apps that seek more permissions from users delay upgrading to the latest version of Android. Among all permissions, the propensity to delay upgrade increases when apps seek permissions to collect user information that are known to be used for ad targeting or are considered to be sneaky (i.e., running in the background). Private user information such as audio conversations (gathered via microphone), location (gathered via GPS), users' frequent connections (gathered via call logs or SMS frequency), user identity (gathered via phone status and identity), user content (gathered via access to phone memory) etc. are often utilized by ad networks to serve targeted advertisements (Goode 2018, Limer 2018). Hence, apps that typically benefit from in-app advertising prefer to retain control over such permissions by delay upgrading.

I also find that the propensity to delay upgrade increases among apps that seek permissions to sensitive information that are non-essential to the app's working. Extant literature on mobile security provides evidence to support this explanation. Displaying

in-app-advertisements when users use those apps is quite a common revenue model among free apps, much more than paid apps (Watanabe et al. 2017). Some apps build revenue models around utilizing users' personal information in ways that users may not agree with (Cimpanu 2018). Some advertisement libraries that apps utilize are known to follow non-standard data policies such as uploading users' sensitive information to remote servers or triggering code from remote servers on apps (Cimpanu 2018, Grace et al. 2012). Such reckless practices by advertisement networks (prone to security and privacy vulnerabilities) are not completely without the apps' knowledge (Grace et al. 2012). My investigation finds that these apps also have incentives to over-seek *non-essential* permissions. Hence, when platforms enact policies to restrict blanket permissions to such sensitive information, such apps would choose to hold off upgrading to this version for as long as possible. My finding is also consistent with prior research that suggests that seeking sensitive information that are relevant to transactions will not be difficult (Zimmer et al. 2010). The findings point towards apps' belief that app users will meaningfully exercise their control over privacy and grant only contextually relevant sensitive information to apps.

In examining the outcomes for apps that delay upgrading to Android 6.0, I find that apps that delay upgrading to Android 6.0 (or higher) significantly reduce over-seeking the *non-essential* permissions. This finding is important because it suggests that giving fine-grained control to users over their private information not only improves apps' data gathering behavior but also that it continues to improve with time. Specifically, this suggests that apps reduce seeking irrelevant permissions that may raise alarms in users' minds. Most importantly, I find delaying upgrade significantly reduces apps'

82

ability to overseek *non-essential* permissions, indicating that early movers among apps find it easier to obtain larger proportion of *non-essential* permissions. One of the reasons for this initial ease of obtaining permissions may be users' uncertainty about breaking app functionality if they deny permissions (Hildenbrand 2018) and such uncertainties would go away by the time late-movers among apps upgrade.

Investigating the marketplace outcomes for apps, I find that delaying upgrade potentially erases most ratings gains that apps would receive upon upgrading. Furthermore, the sub-sample analysis reveals that the fall in the rating are mainly for apps that do not reduce over-seeking *non-essential* permissions. Taken together, the findings suggest that app users are sensitive to apps' over-seeking behavior [32]. I also find that apps' popularity falls upon delaying upgrade. The fall in the popularity of the apps can be attributed to the penalty that Android levies on apps' core quality score for delaying upgrade. These findings imply that apps have a medium- to long-term price to pay if they delay upgrading, despite retaining control over users' information for a bit longer. These findings are quite important in an open hardware platform such as Android. Even if an open platform cannot coerce all apps to upgrade immediately, my findings show that exploiting such flexibility for short-term gains comes at a cost to the apps.

3.5.2 Limitations and Future Research

I acknowledge several limitations of this study. First, similar to prior research on mobile app marketplaces (Garg and Telang 2013, Ghose and Han 2014), my sample

---

[32] Unlike feature improvements which individual apps need to opt-in to, privacy control improvements are applicable to all apps. Therefore, while existing app users won't be unhappy if apps delay feature enhancements, users are likely to be unhappy if they notice that apps are delaying privacy control (a list of all permissions is displayed each time users update the app).

focuses on the most popular apps at a given timeframe (April 2016). The apps have to be installed on an emulator for me to detect when they upgraded, which precludes the use of hundreds of thousands of apps in my analysis. While I put efforts to make my sample as representative as possible and conduct robustness checks that suggest my estimates are conservative, generalizability concerns limit the findings of this study to overall apps, particularly those at the tail-end of the download distribution. Furthermore, my sample is limited to those apps that can be downloaded on emulators from the US Play Store. We need additional studies across a larger sub-sample of apps that would include apps with very number of low downloads as well as those apps that are available only in international markets.

Second, I rely on a statistical technique, combined with textual analysis, to derive the essentiality of permissions to apps' working. While such a technique makes my finding conservative[33], we need more research to measure the essentiality of permissions in alternative ways. Future research can interview app users to understand essentiality from users' perspective. Future work could derive an importance measure of individual apps. Whether and how upgrade decisions vary with an app's relative importance to users would provide additional insights.

Third, while my study shows that apps reduce overseeking *non-essential* permissions, my study does not observe when and how the apps seek such permissions during run-time. There may be some apps that are successful in overseeking *non-essential* permissions by "educating" app users about the necessity of such permissions

---

[33] When *essential* permissions get tagged as *non-essential*, they reduce the effect size of DiD estimates in models estimating effects on *non-essential* permissions. The 75% cut off for *essential* permissions prevents *non-essential* permissions from being categorized as *essential* ones.

for apps' survival. Future research should study app interface designs, specifically the messaging design that may convey to users why certain permissions are necessary to an app's working, thereby avoiding being rejected.

Finally, despite using a reputable matching technique in conjunction with DiD estimation to account for endogeneity of apps' upgrade decisions in a quasi-experimental setting, it is plausible that the self-selected decisions to delay upgrade may be influenced by other unobservable factors. We need more research to leverage exogenous shocks (such as platform level shocks) that would encourage some (or all) apps to upgrade immediately.

Chapter 4: Conclusion and Way Forward

4.1 Conclusion and Contribution to the Literature and Practice

The phenomenal rise of platform economy, which have been around for over a decade now, can be attributed to their success in enabling discovery and transactions between the initial generation of platform users. However, the learnings from managing the earlier users need not translate completely onto the new users, owing to the difference in the sensitivity to issues such as privacy, discrimination, digital nativity and the like. Furthermore, as more and more users flock to platform economy, regulators and policymakers have been involving in platform regulation more frequently now than ever. As a reaction to issues faced by one side of the platform (e.g., discrimination faced by guests on Airbnb) and the accompanying bad press, platforms are introducing policy changes aimed at empowering that side. However, such policy changes could elicit strategic behavior from the other side of the network too. Motivated by a rise in the platform policy changes that reflect both the regulatory mood of the region as well as the changing platform users' preferences and behavior, my dissertation seeks to address the question of *whether* and *how* such policy changes impact all sides of the platform. More importantly, I seek to understand what the unintended consequences of such policy changes are, so that such findings can be used to guide policy decisions.

I choose two distinct and prominent digital platforms: sharing economy platform (Airbnb) and a mobile platform (Android) and two distinct platform policy changes: making screening optional and enabling fine grained control over sensitive information to study the impact of such platform policy changes. In both these contexts, the voluntary nature of such policy adoptions allows me to study *who* adopts such policy

changes, *why* do they adopt such changes and *what* happens to such adopters upon embracing the change. Furthermore, as a study design, I focus on two critical impacts of such policy changes: efficiency and welfare implications, for all sides of the platform (e.g., guests and hosts on Airbnb, developers and users on Android), not just the side that is the intended beneficiary of such policy changes.

The broader contribution of my dissertation is its investigation of what happens when platform design decisions increase the choices of buyers on the platform. Allowing hosts to forgo screening or forcing apps to provide users a fine-grained choice of whether or not to grant individual permissions improves the choices that renters or app users have. Left to market, such a transfer of power (of choice) to consumers is a stretch. Therefore, it is essential for platforms to enact policies that allow for decentralized choices. From my studies, I find that participants may treat cost-benefits differently, but on average, they are perceptive in terms of imposing penalties to actions that do not make sense. In other words, when platforms enact policies that enable individuals to make use of their own decentralized information, I find that there are both efficiency as well as welfare implications.

The first essay informs us that a seemingly counter-intuitive choice of allowing sharing economy participants to forgo screening works favorably to particular sub-group of participants who have traditionally felt marginalized on the platform. The emergent and growing literature on the sharing economy has increasingly attended to the social and economic impacts of online sharing marketplaces (e.g., Einav et al. 2016; Li et al. 2016; Park et al. 2017; Wang et al. 2016; Zervas et al. 2014) and examined how transactions are consummated in such markets (e.g., Cullen and Farronato 2014;

Ye et al. 2018). However, less is known about how market mechanisms should be designed to support the collaborative consumption activities in online sharing markets. Online sharing platforms are not fully regulated, and mostly rely on a variety of self-designed mechanisms to facilitate and govern market transactions (Koopman et al. 2014). Accordingly, they represent an opportunity to critically assess the impacts of market mechanisms and identify those that best serve P2P sharing markets (Luca 2017). My study addresses this research gap, and in doing so, contributes to the literature on market mechanism design and its socio-economic implications, as I elaborate below.

So far, few studies have focused on examining screening mechanisms in online markets from the perspective of the party engaging in screening. In part, this may be because screening mechanisms have been traditionally considered a vital part of P2P online markets to address issues of information asymmetry (e.g., Stiglitz 1975), and indeed, with my study's research context, two-way screening was part of the default market design. Scholars have largely assumed screening is beneficial from the focal agent's perspective, even as they highlight the potential negative impact of screening through increased discrimination from the other agent's perspective, and reduction in market efficiencies through reduced matching between agents (e.g., Cornell and Welch 1996; Fradkin 2015; Pallais 2014; Romanyuk 2016). My study questions the assumption that screening is always beneficial from the focal agent's perspective, and also relaxes the implicit assumption that providers within local areas have homogeneous demand characteristics. Accordingly, my empirical examination of when, and for whom forgoing screening may result in better outcomes reveals that individual agents assess their own relative benefits and costs across implementing and

forgoing screening. In doing so, I provide evidence that the perceived risk of encountering "lemons" in online markets may be offset by greater economic benefits. I show that agents act consistently with this cost-benefit calculus when determining whether forgoing screening is individually optimal given the demand conditions facing their offerings.

The study also contributions to the literature on online markets by highlighting the roles of race/ethnicity and gender in understanding participant behavior. Prior studies have shown the presence of racial discrimination against African American hosts on Airbnb (e.g., Edelman et al. 2017; Kakar et al. 2017), and my results suggest the increased convenience provided through their voluntary forgoing of screening offsets taste-based bias in guests, such that African American hosts benefit from higher occupancy levels at lower costs (in terms of reduced reviewer ratings), relative to White hosts. Similarly, I show that women are more likely to forgo screening than men, and their greater benefits as compared to men largely relate to their ability to maintain price, rather than increase occupancy rates. Regardless of whether through occupancy rates or through listing price, my study provides evidence of higher economic benefits to African American and female hosts when they choose to forgo screening. In doing so, my findings not only complement existing literature on how online market participants respond to mechanism design (e.g., Bapna et al. 2004; Ye et al. 2014), but add to the growing literature on the roles of race/ethnicity and gender in online markets (e.g., Chan and Wang 2017; Edelman et al. 2017).

Practically, my findings provide important implications for online market platform design, and for participants in these markets. Within my focal context, Airbnb has

steadily gained in market share of temporary rentals, and as a recent entrant into the hospitality industry, eroded the dominance of traditional hotel chains (e.g., Zervas et al. 2014). While initially operating under the assumption that two-way screening was critical for information asymmetry reduction, Airbnb's market design precluded the "instant booking" norm in hotel chains. In so doing, it failed to cater to the convenience and time preferences of guests, and also increased concerns regarding discrimination of guests by hosts. My systematic investigation of their recent policy change to enable both traditional and instant booking option provides evidence that guests react positively to instant booking, and this may even offset their taste-based bias. I also show that hosts themselves benefit with having more choices, particularly when they are new entrants, have yet to build their reputation and quality signals, or face lower or declining demand. The option to forgo or implement screening provides them more degrees of freedom in customizing their choices: they can assess their own trade-offs between increased occupancy rate/prices and reduced review ratings, thereby increasing overall benefits in the market. Of course, critical to such outcomes are other features of market design; it may well be that signaling and insurance mechanisms to offset information asymmetry in Airbnb permit screening to be an option.

Moving beyond Airbnb, other online platforms may also benefit from my study, inasmuch as they assess their own overall design, and determine whether forgoing screening can be made optional to serve as an additional sorting tool that differentiates participants with different trade-offs between occupancy, price, and ratings.

My second study on privacy policy change on Android guides the literature on platform mechanism design by showing that altering privacy policies to make choices

less costly to users improves apps' information gathering behaviors. Platforms' proactive measures to reduce users' cost of exercising control over their information may help address the *privacy paradox* that has been documented in privacy literature. By showing that even popular apps systematically lower permission over-seeking for information that are not relevant for apps' working, I highlight the role of platform mechanisms under a power asymmetry in safeguarding users' interest. The fact that users acknowledge the change in privacy control indicates that the Android' strategy to improve the information security, by reducing the costs of making privacy choices, seems successful.

While extant research has investigated the buyers' perspective to information privacy, this study investigates the sellers' responses to exogenous changes in information gathering policies. In doing so, I also add to the literature on early movers by showing that upgrading to the latest version at the earliest is beneficial for an app's outcome in the marketplace, despite potentially disrupting its revenues in the short run. Alternatively, I also show that strategically delaying adhering to policies harm the sellers' outcomes on the platform, not just revenue losses (reduction in *non-essential* permissions and popularity) but also platform reputational losses (erasure of potential gain in ratings). Finally, my research informs platform strategy by suggesting that a fragmented platform such as Android should carefully design its upgrade window. Providing a longer time-horizon may induce strategic behavior among sellers that may not be optimal for the participants or the platform.

In both my studies, I have focused on organically embedding some of the modern text mining and machine learning techniques along with rigorous econometric analysis.

A key strength of modern machine learning/neural network techniques is to uncover the richness of data to guide the model choice and research design. For example, in case of the Airbnb essay, I used neural networks to train a comments analysis model that distinguished positive statements from specific complaints. I also used sentiment analysis techniques to understand how the sentiments have changed with the change in the guest type due to instant booking. I also used neural network based image recognition tool to provide confidence to my manual tagging of host images to determine host race/ethnicity and gender. In the Android App essay, I contribute to the literature on privacy preferences by combining advanced text mining techniques with permission prevalence analysis to distinguish essential sensitive information from non-essential ones. While such a technique is used by platforms such as Google, to detect whether apps are over-seeking permissions or not, this is not too commonly used in social sciences research. By employing a modern vector embedding technique (Word2Vec), I demonstrate that users, despite expressing preferences to controlling sensitive information sharing, are willing to trade-off contextually essential sensitive information to derive utility from apps.

In summary, the studies shed light on critical choices made in market design, and their subsequent impact on market participants' choices and outcomes. I show that rather than a "one size fits all" policy under the assumption that market participants equally desire the same protection against information asymmetry, provision of options to individually assess the benefits and costs associated with policies allows for more decentralized and heterogeneous choices, with associated increases in market outcomes and, thereby, social welfare implications. More importantly, my studies show that

platforms need to take into account the strategic moves by the all sides of the network, while designing and implementing policy changes.

4.2 Way Forward

My dissertation highlights the importance of understanding the consequences of platform policy changes on outcomes for all sides of the platform. Both the studies are observational studies where I studied the outcome based on archival data. A natural extension of my dissertation is to study the impact of mechanism designs in the field by introducing novel mechanism designs on platforms that stand to gain from altering the way users interact on the platform. By carefully analyzing how different sides on such platforms currently interact and identifying the implicit assumptions around such interaction mechanisms, introducing policy changes in a randomized control trial would be useful.

4.2.1 Smart TV Platform

One such platforms would be a traditional content consumption platform such as the Television. Over the past decade, television has transitioned from a broadcast content consumption device to a multi-faceted entertainment device. About 70% of all TVs sold in 2018 are "Smart TVs" (Statista 2018c), i.e., televisions that have internet connectivity to support over-the-air content streaming, internet content consumption or interactive content (for example, Black Mirror: Bandersnatch). Their ability to take input from (and potentially provide output to) multiple connected smart devices and digital assistants has promised to change the way consumers interact with TVs and consume TV content. In essence, the TV has transformed from a simplex (one-sided) communication platform (i.e., from broadcaster to consumer) into an interactive platform. Despite an explosion in over-the-air content consumption, US viewers spend

93

about 79% of their content consumption time consuming linear TV (Nielsen 2018). This has provided Broadcast TV channels and marketers with an opportunity to utilize smart TV platforms for a rich, interactive engagement with consumers.

The current stream of communication between advertisers/marketers and the content consumers on TV has been one-sided; therefore, the ad is broadcast on TV without much customization and micro targeting. The implicit assumption here is that the consumers self-select into remembering the advertisements they are interested in. Therefore, advertisers have relied on creativity to condense advertisement messages and ad-repetition to improve brand recall while at shopping aisles (Elsen et al. 2011; Lehnert et al. 2013). Smart connected TVs are in a perfect position to mitigate the temporal disconnect when the content is consumed and when it is utilized. Technologies on a smart TV can allow TV viewers to "bookmark" an on-TV advertisement and forward the bookmark to their connected smart phone. This can, in turn, help marketers to customize the message as consumers progress through the purchase funnel on a second screen. Figure 4.1 illustrates multi-screen content customization that the smart TV platform enables to advertisers.

The platform level interventions that I propose make use of two unique aspects of linear TV viewing on a smart TV platform. The first is the *time-shift* aspect for an advertisement where a user can postpone the processing of the ad message to a later time (and another device). This means that my interventions can exploit the unique temporal dimension of linear TV viewing (i.e., *when* the TV ad is going to start and end). The second aspect is that this setup provides a multi-level, multi-device interactivity. This means that my interventions can make use of the two-staged (on-TV

and on-Mobile) interactions with self-selected consumers. The former aspect guides my message framing on-TV and the latter aspect guides my intervention design in two-stages.

4.2.1.1 Intervention Message Framing on-TV

The interventions on-TV represent an early stage in a consumer decision-making process. The key goal of on-TV intervention is to persuade viewers, who are in a content-consumption mode, to postpone ad-message processing to their convenient time. To achieve this time-shift, I propose using variants of temporal information as described above. I draw on consumer decision-making literature to frame the persuasive messages. Since the early stage of a consumer's decision-making process has more uncertainty about shopping goals (Lee and Ariely 2006), any intervention that helps reduce this uncertainty increases the response in early stages. Interventions with convenience framing (such as, providing context relevant information (e.g., Goic et al. 2019) or interest customized information (e.g., Fong 2017)), or negative framing (such as scarcity framing (e.g., Luo et al. 2019)) have been shown improve response in early stages of decision-making. Thus, I use convenience and scarcity temporal framing on-TV.

To understand which of these framings (positive vs negative framing) may be more effective, I refer to the risk preference literature. Responding to an on-TV marketing intervention represents a trade-off between convenience (customized product information without remembering and searching) and privacy (exposing their interest to marketers). Thus, for a TV viewer interested in the product, responding to intervention is a risky choice compared to remembering and searching information by

themselves. Kahneman and Tversky (1981) show that people display risk seeking (aversion) when exposed to negative (positive) framing. Therefore, I postulate that negative framing (scarcity framing) better persuades users to respond to on-TV intervention. However, positive framing is shown to be more persuasive under low involvement situations such as with product advertisement (Maheswaran and Meyers-Levy 1990). Thus, it is worth empirically exploring how positive vs negative temporal framing impacts the final outcome (such as download, sales or test drive).

The two-staged design also provides information on other similar viewers' actions, i.e., social information. Theories on social influence suggest that social information cues significantly improve the outcomes for early stages (Aral and Walker 2011; Bakshy et al. 2012) as well as for late stages (Sun et al. 2019). The two distinct components of social influence: social learning and network effect/utility (e.g., Sun et al. 2019) seem to affect consumer decision making at different stages. Social learning cues helps reduce the uncertainty around a product's demand in early stages of decision-making. Similarly, social utility cues help gauge the positive utility from purchasing certain products in late stages of decision-making. For example, cues that indicate how many people have test driven a new car helps gauge the waiting period or the corresponding accessories market. While I expect social information cues to impact both the stages, it is unclear from the literature how much impact social information has on these two stages. Finally, it would also be interesting to study the interaction between social information and the convenience-scarcity framing. Especially since the effectiveness of convenience vs scarcity framing depends on consumers' involvement

in message processing, it would be interesting to understand if social information alters this effectiveness.

4.2.1.2 Intervention Design

The key intuition here is that different types of nudges/cues/information works differently on early vs late stages. As mentioned above, I propose using two variants of temporal information (convenience vs scarcity framing) as on-TV nudges. The social information nudge will be used both on-TV and on-Phone.

4.2.1.2.1 Temporal Information: The convenience framing would be "Bookmark the coupon now, use it at convenience" while the scarcity framing would be "You have 30 seconds to save this coupon". We use these two framing independently and in combination with the social information nudge (this makes a total of four treatment groups and a control group). The idea is to provide short and crisp cues on TV to persuade viewers to time-shift their decision-making to a convenient time.

4.2.1.2.2 Social Information: Other similar users' actions have been shown to be effective in social networking context (e.g., Bakshy et al. 2012). It would be interesting to know *whether* and in what stage (i.e., on-TV vs on-Phone) such information would be effective. An example of a social information would be "3214 viewers watching NPR News in Naperville, IL have bookmarked this coupon".

4.2.2 Policy Simulations

Finally, since we are interested in informing platform policymakers about the impact of such policy changes, another extension of my dissertation would be to simulate policy actions and perform counterfactual analysis of policy enactment. For example, both the policies that I study in my dissertation have looked at the impact of voluntary adoption of the features on their outcomes. A policymaker in Airbnb or Android would

be interested in studying the impact of making such a policy mandatory (instead of optional) or study the impact of altering (shortening or increasing) the time-window of such policy enactment. It would be valuable to setup structural models (e.g., Ghose and Han 2014) and perform a series of policy simulations.

Tables

Table 2.1: Variable Explanation

Variable	Description
Dependent Variables	
$instant\_bookable_{it}$	If listing i has enabled instant booking in month t.
$occupancy_{it}$	Ratio of number of days listing i is booked to number of days the room is marked as available for booking in month t.
$price_{it}$	Price per night stay for listing i in month t (in US dollars).
$rating_{it}$	The cumulative review rating of listing i by the end of month t.
$traditional\_bookable_{it}$	If listing i which switched to instant booking earlier switches back to traditional booking in month t.
Listing Demand Characteristics	
$\Delta occupancy_{it\text{-}1}$	The change in occupancy from month t-2 to month t-1, i.e., $occupancy_{it\text{-}1}$ - $occupancy_{it\text{-}2.}$
$\Delta price_{it\text{-}1}$	The change in price from month t-2 to month t-1, i.e., $price_{it\text{-}1}$ - $price_{it\text{-}2.}$
$\Delta rating_{it\text{-}1}$	The change in rating from month t-2 to month t-1, i.e., $rating_{it\text{-}1}$ - $rating_{it\text{-}2.}$
$occupancy\_change\_percent_{it}$	The percentage change in listing i's monthly occupancy due to instant booking by the end of month t, measured as (average monthly occupancy between the switching month 0 and month t - average monthly occupancy in months before the switching month) / average monthly occupancy in months before the switching month 0.
$price\_change\_percent_{it}$	The percentage change in listing i's price due to instant booking by the end of month t, measured as (average listing price between the switching month 0 and month t - average listing price in months before the switching month / average listing price in months before the switching month 0.
$rating\_change\_percent_{it}$	The percentage change in listing i's review rating due to instant booking by the end of month t, measured as (review rating by the end of month t - review rating before the switching month 0) / review rating before the switching month 0.
$inst\_book\_group_i$	If listing i has ever switched to instant booking.
$post\_switch_t$	If month t is after the month of switching to instant booking by the treated listing (i.e., if t>0).
$switchback\_group_i$	If listing i has even switched to instant booking first and then switched back to traditional booking.
$post\_switchback_t$	If month t is after the month of switching back to traditional booking by the treated listing (i.e., if t>0).
$average\_occupancy_{it}$	The mean of the occupancy of the listing from the start of the panel till the month t.
$competitor\_instant\_bookable_{it}$	Percentage of competitors in listing i's neighborhood that have enabled instant booking in month t.
$competitor\_superhost_{it}$	Percentage of Superhosts in listing i's neighborhood in month t.
$competitor\_review\_count_{it}$	Average number of reviews in month t by listing i's competitors in the same neighborhood.

Host Characteristics

$host\_is\_superhost_{it}$	If the host of listing i is recognized by Airbnb as a Superhost in month t.
$host\_tenure_{it}$	Number of days the host of listing i has been on Airbnb by the end of month t.
$host\_is\_professional_{it}$	If the host of listing i is considered as a professional host (i.e., with 3 or more listings) in month t.[34]
$gender_i$	A vector of dummies indicating if the gender of host of the listing i is male, female or uncertain
$race/ethnicity_i$	A vector of dummies indicating if the race/ethnicity of the host of listing i is African American, Asian, Hispanic, White or uncertain

Control Variables

$review\_count_{it}$	Cumulative number of reviews listing i receives by the end of month t.
min_stay_{it}	Listing i's minimum number of nights required for a booking in month t.
$photos_{it}$	Number of photos posted for listing i in month t.
$popularity_{it}$	Number of times listing i appears in the recommended list of similar listings in month t.
$month\_dummy_t$	A vector of dummies indicating if month t is January, February, etc.
$cancellation\_policy_{it}$	A vector of dummies indicating if the cancellation policy of listing i in month t is flexible, moderate, strict, super-strict or no-refund
$listing\_type_i$	A vector of dummies indicating if the listing i is shared, entire room or entire home/apartment

Table 2.2: Distributions of Host Race/Ethnicity and Host Gender

Race/Ethnicity	Count	Percentage	Gender	Count	Percentage
White	7667	55.73	*female*	6638	48.25
African American	1712	12.44	*male*	5984	43.50
Asian	1369	9.95	*uncertain*	1135	8.25
Hispanic	1367	9.94			
Uncertain	1642	11.94			
Total	13757	100%		13757	100%

---

[34] We follow the definition of professional host by New York State Office of The Attorney General (see https://ag.ny.gov/pdfs/AIRBNB%20REPORT.pdf).

Table 2.3 Switching to Instant Booking

Dependent variable	(1)		(2)	
	Logit - $instant\_bookable_{it}$		Logit - $instant\_bookable_{it}$	
$occupancy_{it-1}$	2.375***	(0.227)	2.423***	(0.228)
$occupancy_{it-1} * occupancy_{it-1}$	-1.436***	(0.202)	-1.432***	(0.203)
$\Delta occupancy_{it-1}$	-0.296***	(0.070)	-0.312***	(0.070)
$price_{it-1}$	-0.001***	(0.225 x10^{-3})	-0.001***	(0.219 x10^{-3})
$\Delta price_{it-1}$	-0.002*	(0.001)	-0.002*	(0.001)
$rating_{it-1}$	-0.010***	(0.001)	-0.009***	(0.001)
$\Delta rating_{it-1}$	0.005	(0.003)	0.005	(0.003)
$review\_count_{it-1}$	0.008***	(0.001)	0.008***	(0.001)
$host\_is\_superhost_{it}$	-0.352***	(0.061)	-0.339***	(0.062)
$host\_tenure_{it}$	-0.001***	(0.690x10^{-4})	-0.001***	(0.695 x10^{-4})
$host\_is\_professional_{it}$	0.695***	(0.041)	0.667***	(0.042)
$male_i$			-0.196***	(0.039)
$uncertain_i$			0.120	(0.063)
$gender_i=female_i$ (baseline)				
$White_i$			-0.137*	(0.061)
$African American_i$			0.311***	(0.070)
$Asian_i$			-0.366***	(0.081)
$Hispanic_i$			0.103	(0.076)
$race/ethnicity_i=uncertain_i$ (baseline)				
min_stay_{it}	-0.019*	(0.009)	-0.020*	(0.009)
$photos_{it}$	0.018***	(0.001)	0.017***	(0.001)
$popularity_{it}$	0.220x10^{-4}	(0.101x10^{-3})	0.302x10^{-4}	(0.105x10^{-3})
$competitor\_instant\_bookable_{it}$	0.881***	(0.116)	0.827***	(0.117)
$competitor\_superhost_{it}$	-0.520**	(0.177)	-0.515**	(0.177)
$competitor\_review\_count_{it}$	-0.001	(0.001)	-0.001	(0.001)
Constant	-3.148***	(0.125)	-3.160***	(0.132)
R-Squared/Pseudo R^2	0.077		0.082	
Cancellation Policy Dummies	YES		YES	
Room Type Control	YES		YES	
Month Dummies	YES		YES	
# of Obs.	117746		117746	
	(Obs. after switch dropped)		(Obs. after switch dropped)	

Table 2.4. Distributions of Host Race/Ethnicity and Host Gender

Race/Ethnicity	Pre-Match	Post-Match (Overall)	Post-Match (Treated)	Post-Match (Control)
White	56.16%	55.10%	52.74%	57.45%
African American	12.21%	13.53%	14.03%	13.02%
Asian	9.78%	9.85%	10.54%	9.17%
Hispanic	9.83%	10.24%	10.43%	10.05%
Uncertain	12.01%	11.28%	12.26%	10.31%
Gender	**Pre-Match**	**Post-Match (Overall)**	**Post-Match (Treated)**	**Post-Match (Control)**
female	48.25%	47.78%	47.42%	48.13%
male	43.50%	44.18%	44.50%	43.85%
uncertain	8.25%	8.04%	8.08%	8.02%

Table 2.5. The Effects of Forgoing Screening

Dependent variable	(1)		(2)		(3)	
	3SLS - $Occupancy_{it}$		3SLS - $Price_{it}$		3SLS - $Rating_{it}$	
$price_{it}$	0.200×10^{-3}*	(0.805×10^{-4})			-0.009***	(0.002)
$rating_{it-1}$	0.001***	(0.137×10^{-3})	0.008	(0.014)	0.097***	(0.003)
$average\_occupancy_{it-1}$	-0.229***	(0.014)	11.458***	(1.472)		
$occupancy_{it}$					13.586***	(1.473)
$review\_count_{it-1}$	0.001***	(0.158×10^{-3})	0.053***	(0.016)	-0.014***	(0.003)
$host\_tenure_{it}$	-0.662×10^{-4}***	(0.176×10^{-4})	0.064***	(0.002)	0.005***	(0.364×10^{-3})
$host\_is\_professional_{it}$	-0.014	(0.020)	5.368**	(2.077)	0.109	(0.426)
$popularity_{it}$	0.235×10^{-4}	(0.195×10^{-4})	-0.006**	(0.002)	-0.220×10^{-3}	(0.409×10^{-4})
$post\_switch_{t}$	0.002	(0.004)	4.716***	(0.447)	0.416***	(0.092)
$inst\_book\_group_{i}$*$post\_switch_{t}$	0.073***	(0.006)	0.319	(0.612)	-1.078***	(0.160)
Constant	0.028	(0.094)	-151.086***	(9.367)	69.593***	(1.957)
Month Dummies	YES		YES		YES	
Listing Fixed Effect	YES		YES		YES	
Cancellation Policy Dummies	YES		YES		YES	
# of Obs.	41,171		41,171		41,171	

*** $p<0.001$, ** $p<0.01$, * $p<0.05$

Table 2.6. The Effects of Forgoing Screening: The Role of Host Race/Ethnicity and Gender

Dependent Variable	(1) 3SLS - $Occupancy_{it}$		(2) 3SLS - $Price_{it}$		(3) 3SLS - $Rating_{it}$	
$price_{it}$	0.199×10^{-3}*	(0.806×10^{-4})			-0.009***	(0.002)
$rating_{it-1}$	0.001***	(0.137×10^{-3})	0.008	(0.014)	0.097***	(0.003)
$average\_occupancy_{it-1}$	-0.231***	(0.014)	11.276***	(1.474)		
$occupancy_{it}$					13.356***	(1.450)
$review\_count_{it-1}$	0.001***	(0.158×10^{-3})	0.054***	(0.016)	-0.013***	(0.003)
$host\_tenure_{it}$	-0.663×10^{-4}***	(0.176×10^{-4})	0.064***	(0.002)	0.005***	(0.361×10^{-3})
$host\_is\_professional_{it}$	-0.012	(0.020)	5.620**	(2.077)	0.049	(0.423)
$popularity_{it}$	0.233×10^{-4}	(0.195×10^{-4})	-0.006**	(0.002)	-0.221×10^{-3}	(0.406×10^{-3})
$post\_switch_t$	0.002	(0.004)	4.708***	(0.446)	0.415***	(0.092)
$inst\_book\_group_i*post\_switch_t$	0.079***	(0.014)	0.078	(1.438)	-1.397***	(0.309)
$inst\_book\_group_i*post\_switch_t*gender_i=male_i$	0.002	(0.010)	-4.280***	(1.007)	-0.017	(0.205)
$inst\_book\_group_i*post\_switch_t*gender_i=uncertain_i$	-0.086***	(0.023)	-1.701	(2.362)	1.569**	(0.491)
$inst\_book\_group_i*post\_switch_t*gender_i=female_i$ (baseline)						
$inst\_book\_group_i*post\_switch_t*ethnicity_i=White_i$	0.004	(0.015)	2.166	(1.496)	0.235	(0.304)
$inst\_book\_group_i*post\_switch_t*ethnicity_i=African American_i$	-0.026	(0.018)	-0.527	(1.843)	1.149**	(0.375)
$inst\_book\_group_i*post\_switch_t*ethnicity_i=Asian_i$	-0.036	(0.020)	7.102***	(2.013)	0.017	(0.413)
$inst\_book\_group_i*post\_switch_t*ethnicity_i=Hispanic_i$	0.025	(0.020)	4.496*	(2.037)	-0.203	(0.416)
$inst\_book\_group_i*post\_switch_t*ethnicity_i=uncertain_i$ (baseline)						
Constant	0.027	(0.094)	-151.445***	(9.364)	69.708***	(1.942)
Month Dummies	YES		YES		YES	
Listing Fixed Effect	YES		YES		YES	
Cancellation Policy Dummies	YES		YES		YES	
# of Obs.	41,171		41,171		41,171	

*** p<0.001, ** p<0.01, * p<0.05

Table 2.7. Comparison of Host Quality: African American vs. White

	Ratings Before Forgoing Screening		
	African American	White	African American vs. White Difference t-test
Switchers	89.76	89.98	-1.47
Non-Switchers	90.80	92.82	-29.07***
Overall	90.52	92.38	-29.23***

	Positive Mentions and Negative Complaints per Comment		
	African American: Before (After)	White: Before (After)	African American vs. White Difference t-test: Before (After)
Positive Mentions	5.47 (4.98)	5.26 (4.72)	1.33 (4.85***)
Room Complaints	0.123 (0.989)	0.119 (0.105)	0.28 (-0.95)
Noise Complaints	0.174 (0.155)	0.157 (0.152)	1.07 (0.38)
Location Complaints	0.000 (0.000)	0.001 (0.001)	-0.83 (-1.75)
Maintenance Complaints	0.032 (0.029)	0.032 (0.030)	-0.03 (-0.08)
Host Complaints	0.079 (0.071)	0.076 (0.076)	0.34 (-0.79)
Amenities Complaints	0.027 (0.035)	0.031 (0.030)	-0.61 (1.54)
Total Negative Complaints	0.435 (0.389)	0.416 (0.393)	0.63 (-0.20)

*** p<0.001, ** p<0.01, * p<0.05.

Table 2.8. Comparison of Host Quality and Price: Female vs. Male

	Ratings Before Forgoing Screening		
	Female	Male	Female vs. Male Difference t-test
Switchers	89.84	89.59	2.11*
Non-Switchers	92.23	92.36	-2.90**
Overall	91.82	91.86	-1.00

	Listing Price at Various Occupancy Ranges		
	Female	Male	Female vs. Male Difference t-test
0-20%	140.45	178.70	-2.61**
20-40%	131.46	151.64	-2.65**
40-60%	145.97	148.27	-0.39
40-80%	145.18	154.35	-2.66**
80-100%	137.30	142.30	-2.73**

*** p<0.001, ** p<0.01, * p<0.05.

Table 2.9. Summary of Major Findings

Research Questions	Major Findings
RQ 1: When and who switches to forgoing screening?	- Forgoing screening is more likely when listings experience: - Mid-range occupancies, or greater recent declines in occupancies. - Listings with higher local competition, and listings whose competitors have forgone screening. - Forgoing screening is more likely by: - New hosts, professional hosts, and non-Superhosts - African American hosts, relative to White hosts. - Female hosts, relative to male hosts.
RQ 2: What are the outcomes associated with forgoing screening, and how do these differ based on host characteristics?	- Occupancy rates increase by an average of 13.52% (1.79 nights per month), listing prices do not change, translating to $300.23 increased revenue per month. Review ratings decrease by an average of 1.20% (1.079 points) - African American hosts have a 3.75% higher increase in occupancy rates, and their review ratings have a 1.28% lower decrease, relative to White hosts. - Female hosts have similar increases in occupancies and decreases in ratings, but without having to decrease listing price by $4.28, as male hosts.

Table 3. 1. Variable Explanation

Variable	Description
Key Variables	
$dangerous\_permissions_{it}$	Count of dangerous permissions sought by app i in month t
$essential\_dangerous\_permissions_{it}$	Count of *essential* dangerous permissions sought by app i in month t
$nonessential\_dangerous\_permissions_{it}$	Count of *non-essential* dangerous permissions sought by app i in month t
$normal\_permissions_{it}$	Count of normal[10] permissions sought by app i in month t
$upgrade_{it}$	A dummy variable indicating if the app i upgrades to Android 6.0 (or above) in the month t.
$rating_{it}$	The cumulative user rating of app i by the end of month t.
$popularity_{it}$	The popularity of app i by the end of month t, computed using the weighted PageRank algorithm
App Characteristics	
$rating\_count_{it}$	Total count of ratings given by raters for app i by the end of month t.
$dayssinceupdate_{it}$	Count of days since updating app i as measured at the end of month t
$screenshots_{it}$	Count of screenshots [35] (including optional video) uploaded by app i to Play Store by the end of month t
$developer\_appcount_{it}$	Number of apps in Play Store by the developer of app i of month t
$filesize_{it}$	File size (in MB) of app i by the end of month t
$inAppAdvertising_{it}$	A dummy variable indicating if app i displays ads in the month t
$inAppPurchases_{it}$	A dummy variable indicating if app i allows in-app purchases in the month t
Control Variables	
$category\_group_i$	App category group that app i belongs to.
$download\_bucket_{it}$	Download bucket (group) that app i belongs to at the end of month t.
$month\_dummy_t$	A vector of dummies that represent if month t is April-2016, May-2017, etc.

Note: Missing values in file sizes (such as "varies with devices" or months when file sizes were not displayed in Play Store) are handled by replacing them with (a) values carried forward from previous months, or (b) mean value of file-size of that app or (c) file-sizes of all apps in a sub-category for the month in the order.

---

[35] Play Store mandates that a minimum of 2 and a maximum of 8 screenshots per target category (phone, tablet, TV and Wear OS) may be uploaded on the Play Store. If the developer has uploaded a video, we count it as a screenshot. https://support.google.com/googleplay/android-developer/answer/1078870?hl=en

Table 3. 2. Analysis of Upgrading to Latest Version of Android

Dependent variable - $upgrade_{it}$	Model (1) Point Estimates		Model (1) Hazard Ratio	
$dangerous\_permissions\_ratio_{it}$	-0.005***	(0.001)	0.995***	(0.001)
$\Delta rating\_count_{it}$	0.000***	(0.000)	1.000***	(0.000)
$rating_{it}$	0.067*	(0.030)	1.069*	(0.032)
$games_i$	-0.136***	(0.036)	0.872***	(0.032)
$personal\_apps_i$	-0.082	(0.043)	0.921	(0.040)
$utility\_mobilespecific_i$	-0.145***	(0.041)	0.865***	(0.036)
$utility\_mobileaccess_i$	0.041	(0.049)	1.041	(0.051)
$learn\_explore_i$	0.082	(0.043)	1.085	(0.047)
$category_i=content\_consumption_i$ (baseline)				
$below\_1million_i$	-0.093***	(0.024)	0.911***	(0.022)
$5million\_10million_i$	0.040	(0.038)	1.040	(0.039)
$10million\_50million_i$	0.182***	(0.043)	1.200***	(0.051)
$above\_50million_i$	0.079	(0.096)	1.082	(0.104)
$download_i=1million\_5million_i$ (baseline)				
$dayssinceupdate_{it}$	-0.004***	(0.000)	0.996***	(0.000)
$screenshots_{it}$	0.012***	(0.002)	1.012***	(0.002)
$filesize_{it}$	0.002***	(0.000)	1.002***	(0.000)
Developer Controls	YES			
Log Likelihood	-71426.18			
AIC	142884.4			
BIC	143043.3			
# of Obs.	152,190			
X^2	1440.90			
P	0.000			

*** $p<0.001$, ** $p<0.01$, * $p<0.05$; Clustered standard errors in parentheses.

Table 3. 3. Analysis of Upgrading to Latest Version of Android

Dependent variable - $upgrade_{it}$	Model (2) *Point Estimates*		Model (2) *Hazard Ratio*	
$location\_coarse_{it}$	0.095**	(0.035)	1.099**	(0.038)
$location\_precise_{it}$	-0.073*	(0.037)	0.930*	(0.034)
$access\_bodysensors_{it}$	0.062	(0.276)	1.064	(0.294)
$make\_phonecall_{it}$	0.022	(0.057)	1.023	(0.058)
$access\_camera_{it}$	0.069*	(0.033)	1.071*	(0.036)
$find\_accounts_{it}$	0.104***	(0.026)	1.110***	(0.029)
$reroute\_outgoingcalls_{it}$	-0.119	(0.116)	0.888	(0.103)
$read\_phonestatus\_and\_identity_{it}$	-0.476***	(0.025)	0.621***	(0.016)
$read\_sms_{it}$	0.027	(0.096)	1.027	(0.098)
$receive\_mms_{it}$	0.053	(0.075)	1.054	(0.079)
$record\_audio_{it}$	-0.107*	(0.042)	0.899*	(0.038)
$send\_sms_{it}$	-0.169	(0.090)	0.845	(0.076)
$make\_sip\_call_{it}$	0.648	(0.334)	1.911	(0.639)
$add\_voicemail_{it}$	-0.198	(0.842)	0.820	(0.691)
$read\_write\_phonememory_{it}$	-0.068*	(0.029)	0.934*	(0.027)
$read\_write\_calendar_{it}$	-0.010	(0.075)	0.990	(0.074)
$read\_write\_calllog_{it}$	-0.637***	(0.097)	0.529***	(0.051)
$read\_write\_contact_{it}$	0.132**	(0.046)	1.141**	(0.052)
$\Delta rating\_count_{it}$	0.000***	(0.000)	1.000***	(0.000)
$rating_{it}$	0.098**	(0.031)	1.103**	(0.034)
$games_i$	-0.107**	(0.037)	0.898**	(0.033)
$personal\_apps_i$	-0.077	(0.044)	0.926	(0.040)
$utility\_mobilespecific_i$	-0.126**	(0.042)	0.882**	(0.037)
$utility\_mobileaccess_i$	0.016	(0.051)	1.016	(0.052)
$learn\_explore_i$	0.028	(0.044)	1.028	(0.045)
$category_i=content\_consumption_i$ (baseline)				
$below\_1million_i$	-0.085***	(0.024)	0.919***	(0.022)
$5million\_10million_i$	0.062	(0.038)	1.063	(0.040)
$10million\_50million_i$	0.221***	(0.043)	1.248***	(0.054)
$above\_50million_i$	0.152	(0.095)	1.164	(0.110)
$download_i=1million\_5million_i$ (baseline)				
$dayssinceupdate_{it}$	-0.004***	(0.000)	0.996***	(0.000)
$screenshots_{it}$	0.011***	(0.002)	1.011***	(0.002)
$filesize_{it}$	0.002***	(0.000)	1.002***	(0.000)
Developer Controls	YES			
Log Likelihood	-71173.32			
AIC	142412.6			
BIC	142740.4			
# of Obs.	152,190			
X^2	1976.79			
P	0.000			

*** p<0.001, ** p<0.01, * p<0.05; Clustered standard errors in parentheses.

Table 3. 4. Analysis of Upgrading to Earlier Version of Android

Dependent variable	(1) Hazard Ratios $upgrade\_api22_{it}$	
$dangerous\_permissions\_ratio_{it}$	1.005***	(0.001)
$\Delta rating\_count_{it}$	1.000***	(0.000)
$rating_{it}$	1.102***	(0.029)
$games_i$	0.798***	(0.026)
$personal\_apps_i$	0.971	(0.036)
$utility\_mobilespecific_i$	0.913*	(0.033)
$utility\_mobileaccess_i$	1.002	(0.043)
$learn\_explore_i$	1.063	(0.040)
$category_i=content\_consumption_i$ (baseline)		
$below\_1million_i$	0.846***	(0.017)
$5million\_10million_i$	1.007	(0.037)
$10million\_50million_i$	1.338***	(0.056)
$above\_50million_i$	1.159	(0.111)
$download_i=1million\_5million_i$ (baseline)		
$dayssinceupdate_{it}$	0.997***	(0.000)
$screenshots_{it}$	1.013***	(0.002)
$filesize_{it}$	1.001***	(0.000)
Developer Controls	YES	
Log Pseudo Likelihood	-84962.99	
AIC	169958	
BIC	170112.6	
# of Obs.	116,157	
X^2	1653.18	
p	0.000	

*** p<0.001, ** p<0.01, * p<0.05; Clustered standard errors in parentheses

Table 3. 5. Analysis of Upgrading to Latest Version of Android

Dependent variable - $upgrade_{it}$	(1) *Hazard Ratios*		(2) *Hazard Ratios*	
essential_dangerous_permissions_ratio$_{it}$	1.005**	(0.002)		
nonessential_dangerous_permissions_ratio$_{it}$			0.993***	(0.001)
Δrating_count$_{it}$	1.000	(0.000)	1.000	(0.000)
rating$_{it}$	1.087**	(0.033)	1.079*	(0.033)
games$_i$	0.898**	(0.033)	0.874***	(0.032)
personal_apps$_i$	0.927	(0.040)	0.924	(0.040)
utility_mobilespecific$_i$	0.867***	(0.036)	0.872***	(0.036)
utility_mobileaccess$_i$	1.000	(0.049)	1.030	(0.050)
learn_explore$_i$	1.090*	(0.047)	1.079	(0.046)
category$_i$=content_consumption$_i$ (baseline)				
below_1million$_i$	0.916***	(0.022)	0.912***	(0.022)
5million_10million$_i$	1.043	(0.040)	1.049	(0.040)
10million_50million$_i$	1.214***	(0.052)	1.221***	(0.053)
above_50million$_i$	1.317*	(0.146)	1.334**	(0.148)
download$_i$=1million_5million$_i$ (baseline)				
dayssinceupdate$_{it}$	0.996***	(0.000)	0.996***	(0.000)
screenshots$_{it}$	1.012***	(0.002)	1.012***	(0.002)
filesize$_{it}$	1.002***	(0.000)	1.002***	(0.000)
Developer Controls	YES		YES	
Log Likelihood	-71444.006		-71421.54	
AIC	142920		142875.1	
BIC	143078.9		143034	
# of Obs.	152,190		152,190	
X^2	1347.76		1393.79	
P	0.000		0.000	

*** p<0.001, ** p<0.01, * p<0.05; Clustered standard errors in parentheses.

Table 3. 6. 2SRI with Individual Frailty (2SRI-F) approach
Panel A: First Stage Linear Models – essential and non-essential dangerous permissions

Dependent variable	(1) Linear FE Model (First Stage) essential permissions		(2) Linear FE Model (First Stage) non-essential permissions	
$inAppAdvertising_{it}$	-0.026	(0.027)	0.128**	(0.042)
$\Delta rating\_count_{it}$	-0.000***	(0.000)	-0.000**	(0.000)
$rating_{it}$	-0.022	(0.176)	0.562*	(0.280)
$below\_1million_i$	-0.128***	(0.038)	-0.075	(0.061)
$5million\_10million_i$	0.072	(0.047)	0.208**	(0.075)
$10million\_50million_i$	-0.061	(0.077)	0.078	(0.123)
$above\_50million_i$	-0.316*	(0.146)	0.635**	(0.232)
$download_i=1million\_5million_i$ (baseline)				
$dayssinceupdate_{it}$	0.000***	(0.000)	0.004***	(0.000)
$screenshots_{it}$	0.009	(0.017)	-0.062*	(0.027)
$filesize_{it}$	0.002**	(0.000)	0.002*	(0.001)
$Constant$	13.503***	(0.752)	21.010***	(1.197)
Developer Controls	YES		YES	
# of Obs.	152,190		152,190	
R-Squared	0.813		0.867	

*** p<0.001, ** p<0.01, * p<0.05;

Panel B: Proportional Hazard – essential and non-essential dangerous Permissions

Dependent variable	(1) CoxPH Model (Second Stage) essential Permissions		(2) CoxPH Model (Second Stage) non-essential Permissions	
$nonessential\_dangerous\_permissions\_ratio_{it}$			-0.008***	(0.001)
$essential\_dangerous\_permissions\_ratio_{it}$	0.007***	0.007***		
$\Delta rating\_count_{it}$	0.000***	0.000***	0.000***	(0.000)
$rating_{it}$	0.075*	0.075*	0.066*	(0.032)
$games_i$	-0.108 ***	-0.108 ***	-0.139***	(0.039)
$personal\_apps_i$	-0.082	-0.082	-0.081	(0.046)
$utility\_mobilespecific_i$	-0.149***	-0.149***	-0.140 ***	(0.044)
$utility\_mobileaccess_i$	-0.005	-0.005	0.032	(0.054)
$learn\_explore_i$	0.086+	0.086+	0.078	(0.047)
$category_i=content\_consumption_i$ (baseline)				
$below\_1million_i$	-0.086 ***	-0.086 ***	-0.092***	(0.026)
$5million\_10million_i$	0.037	0.037	0.042	(0.042)
$10million\_50million_i$	0.176***	0.176***	0.182***	(0.041)
$above\_50million_i$	0.084	0.084	0.085	(0.084)
$download_i=1million\_5million_i$ (baseline)				
$dayssinceupdate_{it}$	-0.004***	-0.004***	-0.004***	(0.000)
$screenshots_{it}$	0.012***	0.012***	0.012***	(0.002)
$filesize_{it}$	0.002***	0.002***	0.002***	(0.000)
$residuals\_first\_stage_{it}$	-0.007	-0.007	0.002	(0.002)
Developer Controls	YES		YES	
Log Likelihood	-71424.4		-71403.8	
AIC	142903.238		142862.037	
# of Obs.	152,190		152,190	
X^2	1922(p=0.000)		1963 (p=0.000)	

*** p<0.001, ** p<0.01, * p<0.05;

Estimation induced random error is accounted by using individual frailties in the Second Stage.

F-test for instrument in first stage for model (1) is 0.99, indicating that the instrument is weak. Kleibergen-Paap rk LM statistic cannot reject (p val. =0.321) the null hypothesis that first stage is under-identified for model (1).

F-test for instrument in first stage for model (2) is 9.093, which is greater than Stock-Yogo F test value at 15% maximal IV size (i.e., 8.96). Kleibergen-Paap rk LM statistic rejects (p val. =0.003) the null hypothesis that first stage is under-identified for model (2).

Table 3. 7. Covariate Balance

Covariates	Pre-Match (Overall)	Pre-Match (Treated)	Pre-Match (Control)	Difference	t-test	Post-Match (Treated)	Post-Match (Control)	Difference	t-test
Rating	4.01	4.03	4.00	-0.03	-17.53	4.00	4.01	0.01	1.20
Rating Count	123188	153233	106024	-47209	-12.06	73624	75736	2112	0.25
File Size	24.00	25.28	23.30	-1.98	-17.19	23.79	24.45	0.66	1.22
Screenshots	12.26	12.55	12.09	-0.46	-17.86	12.73	12.51	-0.22	-1.45
dayssinceupdate	193.88	122.57	234.63	112.06	122.40	170.68	179.30	8.63	1.74
inAppAdvertising	0.17	0.182	0.164	-0.019	-10.68	0.148	0.145	-0.003	-0.27
inAppPurchases	0.30	0.330	0.285	0.045	-20.92	0.339	0.334	0.006	-0.44

Note: Apps that were already in Android 6.0 at the start of the panel are excluded from PSM DiD analysis.

Table 3. 8. Analysis of Effects of Delaying Upgrading to Latest Version
Panel A: Impact on *essential* permissions

Dependent Variable	PSM DID - *essential_ dangerous_permissions$_{it}$*				PSM DID - *essential_ dangerous_permissions$_{it}$*			
	(1)		(2)		(3)		(4)	
Δrating_count$_{it}$	0.000	(0.000)	0.000	(0.000)	0.000	(0.000)	0.000	(0.000)
rating$_{it}$	0.154***	(0.039)	0.259***	(0.048)	0.154***	(0.039)	0.263***	(0.048)
post_upgrade$_t$	0.019***	(0.004)	0.019***	(0.005)	0.019***	(0.004)	0.018***	(0.005)
upgrade_group$_i$*post_upgrade$_t$	-0.012	(0.006)	-0.008	(0.006)	-0.019	(0.012)	-0.045***	(0.012)
late_upgrade_group$_i$*post_upgrade$_t$					0.010	(0.012)	0.051***	(0.012)
dayssinceupdate$_{it}$	0.000***	(0.000)	-0.000*	(0.000)	0.000***	(0.000)	-0.000	(0.000)
screenshots$_{it}$	0.011***	(0.003)	0.012***	(0.003)	0.011***	(0.003)	0.012***	(0.003)
filesize$_{it}$	0.000	(0.000)	-0.000	(0.000)	0.000	(0.000)	-0.000	(0.000)
developer_appcount$_{it}$	-0.000***	(0.000)	-0.000	(0.000)	-0.000***	(0.000)	-0.000	(0.000)
Constant	0.431**	(0.166)	-0.239	(0.198)	0.429**	(0.166)	-0.260	(0.199)
App Fixed Effect	YES		YES		YES		YES	
Month Dummies	YES		YES		YES		YES	
Download Bucket Dummies	YES		YES		YES		YES	
Weighted by population	NO		YES		NO		YES	
# of Obs.	106065		106065		106065		106065	
R-Squared	0. 139		0.108		0.140		0.109	

*** p<0.001, ** p<0.01, * p<0.05;

Heteroskedastic Robust Standard Errors in Parenthesis

	Panel B: Impact on *nonessential* permissions							
Dependent Variable	PSM DID - *nonessential_ dangerous_permissions$_{it}$*				PSM DID - *nonessential_ dangerous_permission*			
	(1)		(2)		(3)		(4)	
Δ*rating_count$_{it}$*	-0.000***	(0.000)	-0.000***	(0.000)	-0.000***	(0.000)	-0.000***	(0.000
rating$_{it}$	-0.258**	(0.100)	-0.380***	(0.102)	-0.263**	(0.100)	-0.389***	(0.102
post_upgrade$_t$	-0.011	(0.008)	0.009	(0.009)	-0.009	(0.008)	0.012	(0.009
upgrade_group$_i$*post_upgrade$_t$	-0.162***	(0.013)	-0.176***	(0.014)	-0.101***	(0.023)	-0.110***	(0.024
late_upgrade_group$_i$*post_upgrade$_t$					-0.083***	(0.025)	-0.092***	(0.026
dayssinceupdate$_{it}$	0.000***	(0.000)	0.000***	(0.000)	0.000***	(0.000)	0.000***	(0.000
screenshots$_{it}$	-0.012	(0.007)	-0.018*	(0.008)	-0.012	(0.007)	-0.019*	(0.008
filesize$_{it}$	0.001***	(0.000)	0.002***	(0.001)	0.001***	(0.000)	0.002***	(0.001
developer_appcount$_{it}$	0.004***	(0.000)	0.004***	(0.000)	0.004***	(0.000)	0.004***	(0.000
Constant	3.719***	(0.425)	4.114***	(0.435)	3.735***	(0.425)	4.150***	(0.435
App Fixed Effect	YES		YES		YES		YES	
Month Dummies	YES		YES		YES		YES	
Download Bucket Dummies	YES		YES		YES		YES	
Weighted by population	NO		YES		NO		YES	
# of Obs.	106065		106065		106065		106065	
R-Squared	0.025		0.026		0.025		0.026	

*** p<0.001, ** p<0.01, * p<0.05;
  Heteroskedastic Robust Standard Errors in Parenthesis

Panel C: Impact on Apps' Ratings

Dependent Variable	PSM DID – log(rating$_{it}$)				PSM DID – log(rating$_{it}$)			
	(1)		(2)		(3)		(4)	
dangerous_permissions_ratio$_{it}$	0.000	(0.001)	0.003***	(0.001)	0.000	(0.001)	0.003***	(0.001)
Δrating_count$_{it}$	0.000***	(0.000)	0.000***	(0.000)	0.000***	(0.000)	0.000***	(0.000)
post_upgrade$_t$	-0.090***	(0.011)	-0.077***	(0.010)	-0.087***	(0.011)	-0.073***	(0.010)
upgrade_group$_i$*post_upgrade$_t$	0.100***	(0.016)	0.109***	(0.015)	0.168***	(0.033)	0.203***	(0.030)
late_upgrade_group$_i$*post_upgrade$_t$					-0.092**	(0.034)	-0.129***	(0.031)
dayssinceupdate$_{it}$	-0.000	(0.000)	-0.000***	(0.000)	-0.000	(0.000)	-0.000***	(0.000)
screenshots$_{it}$	0.014	(0.008)	0.025	(0.013)	0.013	(0.007)	0.025	(0.013)
filesize$_{it}$	0.000	(0.000)	-0.000	(0.000)	0.000	(0.000)	-0.000	(0.000)
developer_appcount$_{it}$	0.003***	(0.000)	0.003***	(0.000)	0.003***	(0.000)	0.003***	(0.000)
Constant	139.722***	(0.133)	137.904***	(0.206)	139.719***	(0.133)	137.907***	(0.206)
App Fixed Effect	YES		YES		YES		YES	
Month Dummies	YES		YES		YES		YES	
Download Bucket Dummies	YES		YES		YES		YES	
Weighted by population	NO		YES		NO		YES	
# of Obs.	106065		106065		106065		106065	
R-Squared	0.107		0.175		0.107		0.176	

*** p<0.001, ** p<0.01, * p<0.05;

Heteroskedastic Robust Standard Errors in Parenthesis

Panel D: Impact on Apps' Popularity

| Dependent Variable	PSM DID – popularity$_{it}$				PSM DID - popularity$_{it}$			
	(1)		(2)		(3)		(4)	
dangerous_permissions_ratio$_{it}$	0.044***	(0.006)	0.033***	(0.007)	0.044***	(0.006)	0.033***	(0.007)
Δrating_count$_{it}$	-0.000	(0.000)	-0.000*	(0.000)	-0.000	(0.000)	-0.000*	(0.000)
rating$_{it}$	14.038***	(1.015)	11.805***	(1.283)	13.992***	(1.014)	11.753***	(1.283)
post_upgrade$_t$	-0.036	(0.103)	-0.142	(0.107)	-0.010	(0.103)	-0.124	(0.107)
upgrade_group$_i$*post_upgrade$_t$	-0.478***	(0.130)	-0.562***	(0.137)	0.126	(0.211)	-0.155	(0.213)
late_upgrade_group$_i$*post_upgrade$_t$					-0.822***	(0.223)	-0.561*	(0.230)
dayssinceupdate$_{it}$	-0.005***	(0.000)	-0.005***	(0.000)	-0.005***	(0.000)	-0.005***	(0.000)
screenshots$_{it}$	-0.309*	(0.155)	-0.042	(0.179)	-0.310*	(0.154)	-0.044	(0.179)
filesize$_{it}$	0.006*	(0.003)	0.013***	(0.003)	0.006*	(0.003)	0.013***	(0.003)
developer_appcount$_{it}$	-0.042***	(0.004)	-0.016***	(0.002)	-0.042***	(0.004)	-0.016***	(0.002)
Constant	-33.563***	(5.017)	-32.280***	(5.884)	-33.401***	(5.012)	-32.061***	(5.881)
App Fixed Effect	YES		YES		YES		YES	
Month Dummies	YES		YES		YES		YES	
Download Bucket Dummies	YES		YES		YES		YES	
Weighted by population	NO		YES		NO		YES	
# of Obs.	106065		106065		106065		106065	
R-Squared	0.038		0.028		0.038		0.028	

*** p<0.001, ** p<0.01, * p<0.05;

Heteroskedastic Robust Standard Errors in Parenthesis

Figures
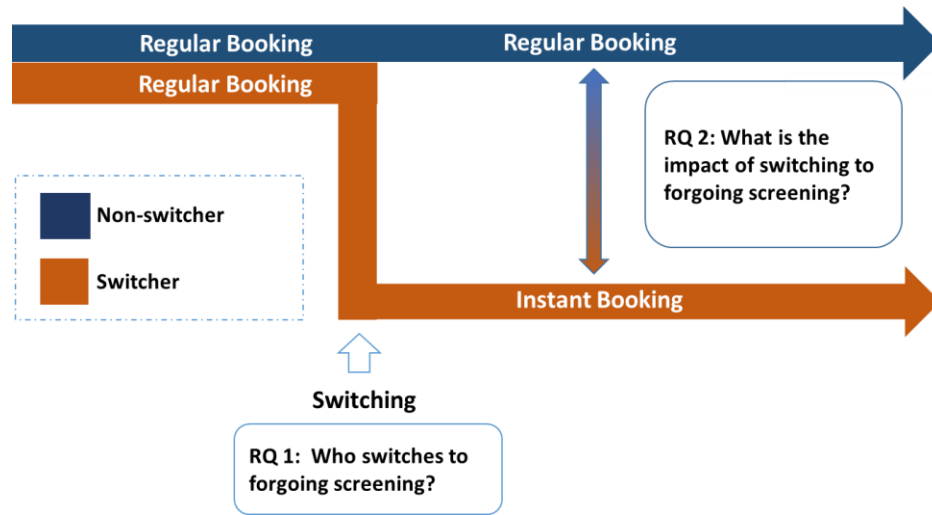
Figure 2. 1. Research Framework
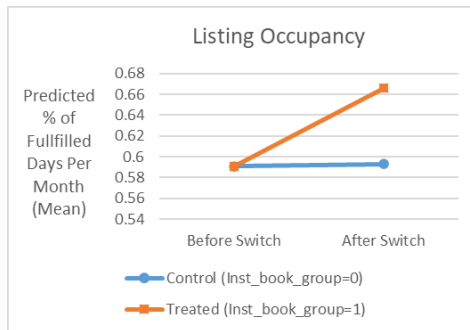
2.2(a): The Effect on Listing Occupancy

2.2(b): The Effect on Listing Price

2.2(c): The Effect on Listing Rating

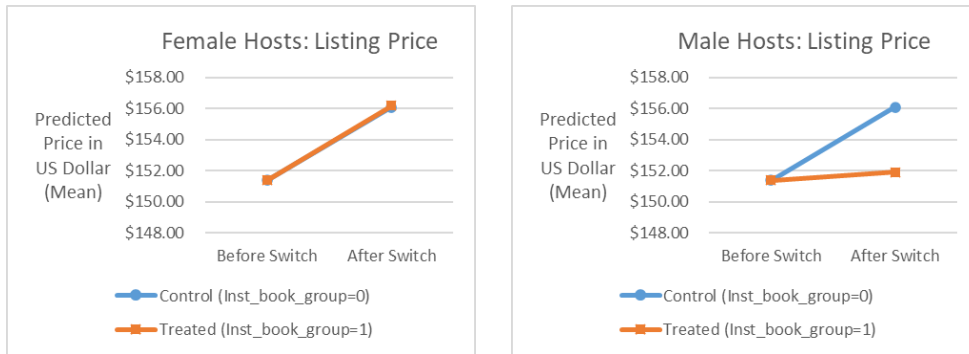**Figure 2. 2. Impact of Forgoing screening on the three variables of interest**

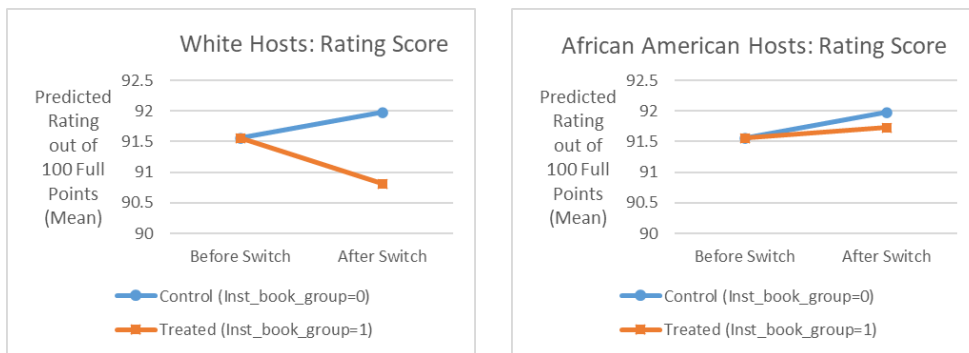Figure 2.3(a): The Effect on Listing Price: Female vs. Male Hosts

Figure 2.3(b): The Effect on Listing Rating: White vs. African American Hosts

Figure 2. 3. Heterogeneity in Outcomes of Forgoing Screening
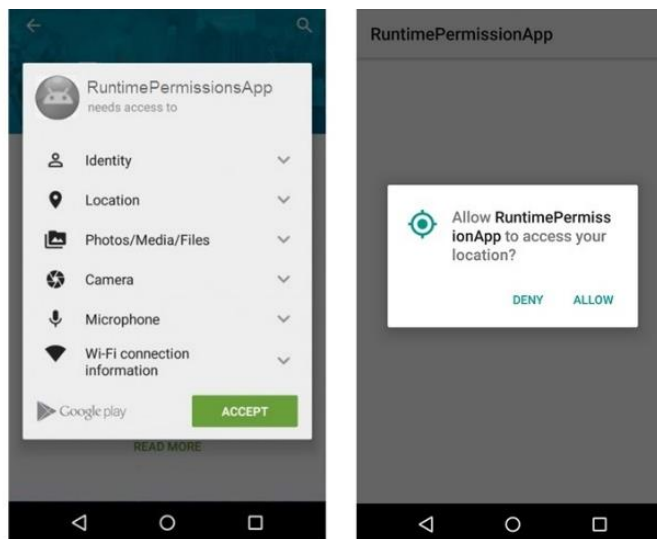
Figure 3. 1. Change in permission-seeking from download to run-time
(Photo Courtesy: GeneXus Community)

Document editing apps

Calendar related apps

Keyboard related apps

Phone utility and information apps

Fig 3.2(A): Distribution of words in sub-categories created from category "Productivity"

Fashion related apps

Messaging and photo apps

Buy and Sell information app

On-Demand order apps

Fig 3.2(B): Distribution of words in sub-categories created from category "Lifestyle"

Figure 3. 2. Sub-Categorization using Vector Embedding Technique

Figure 3. 3. List of "similar apps" from Play Store
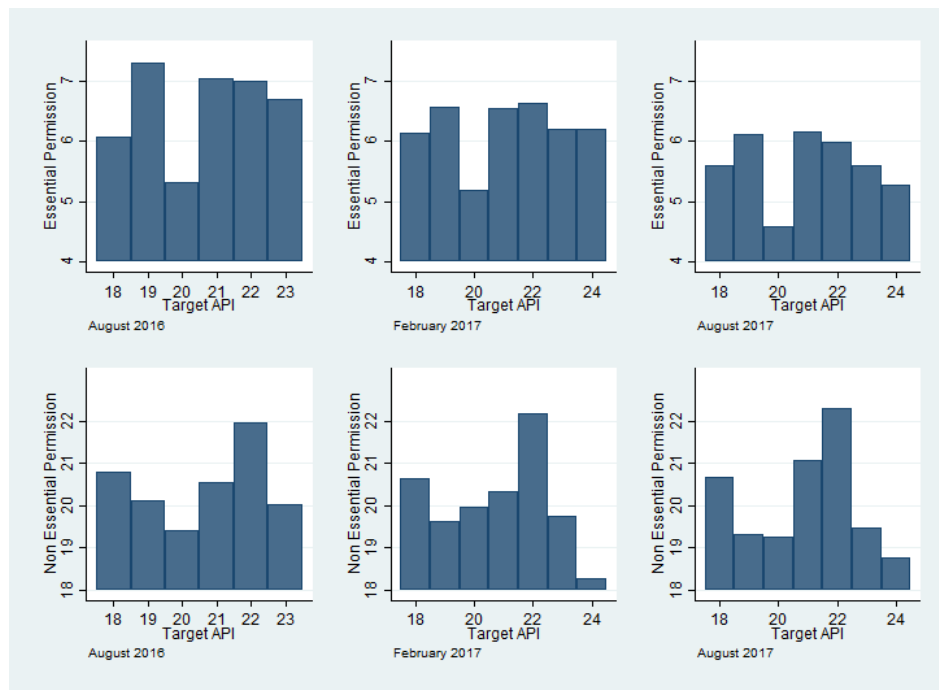
Figure 3. 4. Essential and non-essential permissions over time and targetAPI
Note: API 20 is Watch OS version, hence the deviation.

Figure 3. 5. Parallel Trends Assumption
Note: Time 0 is the month of matching

Figure 4. 1. Multiple levels of interaction on multiple screens
(Photo Courtesy: Sensara Inc.)

Appendices
Table A1. Summary Statistics

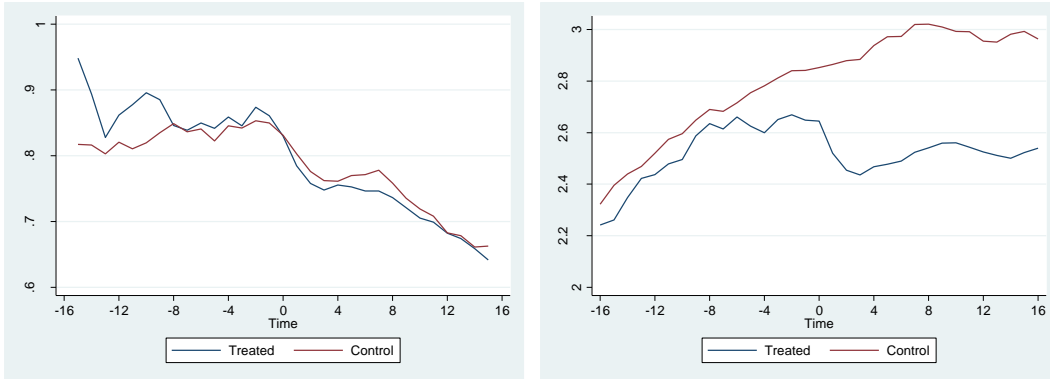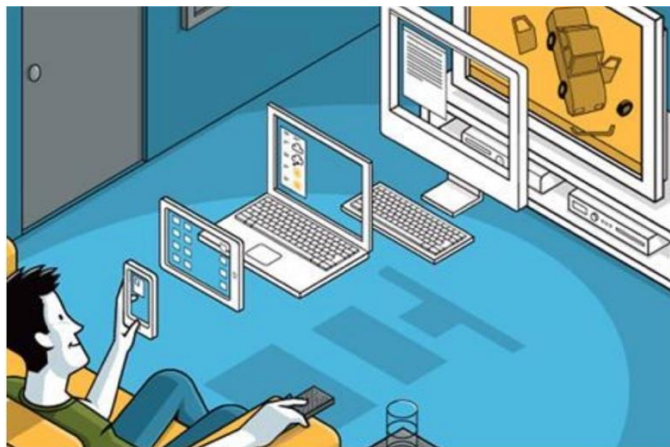| Variables	# of Obs.	Mean	Std. Dev.	Min	Max
Dependent Variables					
$instant\_bookable_{it}$	196155	0.16	0.36	0	1
$occupancy_{it}$	193184	0.54	0.38	0	1
$price_{it}$	196155	167.87	122.26	10	1000
$rating_{it}$	194535	89.55	14.08	20	100
$traditional\_bookable_{it}$	46867	0.34	0.48	0	1
Listing Demand Characteristics					
$inst\_book\_group_i$	196155	0.31	0.46	0	1
$post\_switch_t$	196155	0.18	0.38	0	1
$switchback\_group_i$	196155	0.19	0.39	0	1
$post\_switchback_t$	196155	0.10	0.30	0	1
$average\_occupancy_{it}$	196155	0.53	0.28	0	1
$competitor\_instant\_bookable_{it}$	196155	0.19	0.15	0	1
$competitor\_superhost_{it}$	172900	0.10	0.11	0	1
$competitor\_review\_count_{it}$	196155	35.22	35.10	0	263.42
Host Characteristics					
$host\_is\_superhost_{it}$	172908	0.12	0.32	0	1
$host\_tenure_{it}$	196155	751.08	520.18	3	3095
$host\_is\_professional_{it}$	196155	0.19	0.39	0	1
Control Variables					
$review\_count_{it}$	196152	32.18	36.28	0	408
min_stay_{it}	195869	2.56	3.25	1	180
$photos_{it}$	196152	16.05	11.36	1	197
$popularity_{it}$	196155	11.90	56.93	0	11583

Note: (a) These statistics are based on the sample before propensity score matching.

(b) $rating_{it}$ captures the guests' overall satisfaction score of the listing, which is a more accurate and nuanced measure of a listing's review rating than the approximated star rating (1 to 5)

	Variables	# Obs.	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)	(19)	(20)
(1)	instant_bookable$_{it}$	196155	1.00																			
(2)	occupancy$_{it}$	193184	**0.16**	1.00																		
(3)	price$_{it}$	196155	**-0.09**	**-0.07**	1.00																	
(4)	rating$_{it}$	194535	**-0.01**	**0.20**	**0.02**	1.00																
(5)	traditional_bookable$_{it}$	46867	**-0.73**	**-0.11**	**0.14**	**0.10**	1.00															
(6)	inst_book_group$_i$	196155	**0.65**	**0.15**	**-0.09**	**0.01**	.	1.00														
(7)	post_switch$_t$	196155	**0.48**	**0.10**	**-0.03**	**0.03**	**0.60**	**0.69**	1.00													
(8)	switchback_group$_i$	196155	**0.27**	**0.07**	**-0.05**	**0.02**	**0.39**	**0.72**	**0.73**	1.00												
(9)	post_switchback$_t$	196155	**0.03**	**0.05**	0.00	**0.03**	**0.74**	**0.50**	**0.68**	**0.70**	1.00											
(10)	average_occupancy$_{it}$	196155	**0.18**	**0.73**	**-0.11**	**0.26**	**-0.12**	**0.17**	**0.10**	**0.08**	**0.06**	1.00										
(11)	competitor_instant_bookable$_{it}$	196155	**0.12**	**0.06**	**-0.17**	**0.05**	**-0.02**	**0.11**	**0.08**	**0.04**	**0.07**	**0.14**	1.00									
(12)	competitor_superhost$_{it}$	172900	**0.03**	**0.12**	**0.09**	**0.13**	**0.07**	**0.03**	**0.06**	**0.02**	**0.06**	**0.19**	**0.13**	1.00								
(13)	competitor_review_count$_{it}$	196155	**0.11**	**0.31**	**-0.02**	**0.14**	0.00	**0.08**	**0.07**	**0.01**	**0.07**	**0.49**	**0.23**	**0.27**	1.00							
(14)	host_is_superhost$_{it}$	172908	**-0.03**	**0.11**	**0.05**	**0.18**	**0.06**	**-0.03**	0.00	**-0.01**	0.00	**0.15**	**0.05**	**0.15**	**0.12**	1.00						
(15)	host_tenure$_{it}$	196155	**-0.08**	0.00	**0.08**	**0.05**	**0.15**	**-0.09**	**-0.02**	**-0.06**	0.01	**0.09**	**0.09**	**0.18**	**0.49**	**0.04**	1.00					
(16)	host_is_professional$_{it}$	196155	**0.16**	**0.02**	**-0.18**	**-0.07**	**-0.11**	**0.17**	**0.07**	**0.08**	**0.04**	**0.03**	**0.10**	**-0.03**	**-0.01**	**-0.09**	**-0.02**	1.00				
(17)	review_count$_{it}$	196152	**0.15**	**0.35**	**-0.06**	**0.14**	**-0.06**	**0.12**	**0.08**	**0.02**	**0.06**	**0.54**	**0.21**	**0.24**	**0.83**	**0.13**	**0.48**	**0.03**	1.00			
(18)	min_stay$_{it}$	195869	**-0.04**	**-0.03**	**0.06**	**0.01**	**0.08**	**-0.03**	0.00	**-0.01**	0.01	**0.01**	**-0.02**	0.00	**0.01**	**0.01**	**0.12**	**-0.03**	**-0.04**	1.00		
(19)	photos$_{it}$	196152	**0.08**	**0.11**	**0.17**	**0.13**	**-0.01**	**0.09**	**0.06**	**0.06**	**0.05**	**0.15**	**0.03**	**0.11**	**0.17**	**0.10**	**0.13**	**0.09**	**0.19**	**0.02**	1.00	
(20)	popularity$_{it}$	196155	**0.01**	**0.02**	**-0.04**	**0.01**	**-0.01**	**0.01**	**0.00**	0.00	0.00	**0.02**	**0.01**	0.00	**0.02**	0.00	**0.06**	**0.01**	**0.04**	**-0.01**	0.00	1.00

Bold: correlations significant at p<0.05 level (based on all samples before propensity score matching)

Table A3. Sub-Sample Analysis of Switching to Instant Booking

Dependent variable	Female Sub-Sample		Male Sub-Sample	
	(1)		(2)	
	Logit - *instant_bookable$_{it}$*		Logit - *instant_bookable$_{it}$*	
occupancy$_{it-1}$	2.856***	(0.317)	2.021***	(0.368)
*occupancy$_{it-1}$ * occupancy$_{it-1}$*	-1.838***	(0.281)	-1.211***	(0.328)
Δoccupancy$_{it-1}$	-0.250**	(0.095)	-0.267*	(0.116)
price$_{it-1}$	0.348 x10^{-3}	(0.283 x10^{-3})	-0.002***	(0.394 x10^{-3})
Δprice$_{it-1}$	-0.002	(0.001)	-0.004**	(0.001)
rating$_{it-1}$	-0.008***	(0.002)	-0.010***	(0.002)
Δrating$_{it-1}$	0.006	(0.005)	0.003	(0.005)
review_count$_{it-1}$	0.008***	(0.001)	0.008***	(0.001)
host_is_superhost$_{it}$	-0.345***	(0.088)	-0.268**	(0.098)
host_tenure$_{it}$	-0.001***	(0.883 x10^{-4})	-0.001***	(0.118 x10^{-3})
host_is_professional$_{it}$	0.651***	(0.060)	0.583***	(0.073)
African American$_i$	0.708***	(0.068)	0.226*	(0.091)
Asian$_i$	-0.166	(0.091)	-0.099	(0.102)
Hispanic$_i$	0.313***	(0.078)	-0.176	(0.104)
Others$_i$	0.291**	(0.096)	0.111	(0.100)
White$_i$ (baseline)				
min_stay$_{it}$	-0.040**	(0.014)	0.008	(0.009)
photos$_{it}$	0.027***	(0.002)	0.004	(0.003)
popularity$_{it}$	-0.187 x10^{-3}	(0.311 x10^{-3})	0.197 x10^{-4}	(0.955 x10^{-4})
competitor_instant_bookable$_{it}$	0.706***	(0.172)	0.654***	(0.186)
competitor_superhost$_{it}$	-0.245	(0.240)	-0.481	(0.303)
competitor_review_count$_{it}$	-0.003*	(0.001)	-0.001	(0.001)
Constant	-3.490***	(0.188)	-2.520***	(0.214)
Pseudo R-Squared	0.081		0.070	
Cancellation Policy Dummies	YES		YES	
Room Type Control	YES		YES	
Month Dummies	YES		YES	
# of Obs.	56705		51634	

Table A4. Validation of Parallel Trend Assumption for DiD Estimation

Dependent Variable	(1) 3SLS - $Occupancy_{it}$		(2) 3SLS - $Price_{it}$		(3) 3SLS - $Rating_{it}$	
$price_{it}$	0.538×10^{-3}***	(0.994×10^{-4})			-0.020***	(0.002)
$rating_{it-1}$	-0.001***	(0.165×10^{-3})	-0.116***	(0.013)	0.091***	(0.003)
$average\ occupancy_{it-1}$	-0.242***	(0.019)	23.784***	(1.517)		
$occupancy_{it}$					-3.768***	(0.328)
$review\_count_{it-1}$	-0.001**	(0.346×10^{-3})	-0.103***	(0.025)	-0.038***	(0.005)
$host\_tenure_{it}$	0.232×10^{-3}***	(0.240×10^{-4})	0.025***	(0.002)	0.006***	(0.450×10^{-3})
$host\_is\_professional_{it}$	-0.006	(0.032)	1.389	(2.414)	-0.243	(0.522)
$popularity_{it}$	-0.171×10^{-3}*	(0.696×10^{-4})	-0.023***	(0.005)	-0.349×10^{-3}	(0.001)
$inst\_book\_group_i * time_{t-11}$	-0.043	(0.063)	-4.098	(5.077)	-2.345*	(1.098)
$inst\_book\_group_i * time_{t-10}$	0.076	(0.060)	-3.836	(4.788)	-1.790^{+}	(1.035)
$inst\_book\_group_i * time_{t-9}$	-0.008	(0.057)	-2.946	(4.573)	-0.184	(0.988)
$inst\_book\_group_i * time_{t-8}$	-0.026	(0.056)	-1.080	(4.464)	-0.741	(0.965)
$inst\_book\_group_i * time_{t-7}$	0.023	(0.055)	-2.339	(4.379)	-0.805	(0.947)
$inst\_book\_group_i * time_{t-6}$	0.036	(0.054)	-2.110	(4.322)	-0.728	(0.934)
$inst\_book\_group_i * time_{t-5}$	0.041	(0.053)	-1.234	(4.286)	-1.145	(0.926)
$inst\_book\_group_i * time_{t-4}$	0.020	(0.053)	-1.630	(4.256)	-1.063	(0.920)
$inst\_book\_group_i * time_{t-3}$	0.015	(0.053)	-0.420	(4.227)	-0.853	(0.914)
$inst\_book\_group_i * time_{t-2}$	0.026	(0.053)	0.009	(4.212)	-1.212	(0.910)
$inst\_book\_group_i * time_{t-1}$	0.027	(0.052)	1.199	(4.196)	-0.930	(0.907)
$inst\_book\_group_i * time_t$			Omitted Baseline			
$inst\_book\_group_i * time_{t+1}$	0.076	(0.052)	0.552	(4.191)	-0.855	(0.906)
$inst\_book\_group_i * time_{t+2}$	0.080	(0.052)	1.162	(4.194)	-0.810	(0.907)
$inst\_book\_group_i * time_{t+3}$	0.084	(0.052)	2.200	(4.194)	-0.745	(0.907)
$inst\_book\_group_i * time_{t+4}$	0.090^{+}	(0.052)	3.661	(4.197)	-0.677	(0.908)
$inst\_book\_group_i * time_{t+5}$	0.072	(0.052)	2.484	(4.206)	-0.728	(0.909)
$Constant$	-0.507***	(0.115)	-22.298**	(7.789)	60.607***	(2.090)
Month Dummies	YES		YES		YES	
Listing Fixed Effect	YES		YES		YES	
Cancellation Policy Dummies	YES		YES		YES	
# of Obs.	21,784		21,784		21,784	

*** $p<0.001$, ** $p<0.01$, * $p<0.05$; $^{+}$ $p<0.1$

Testing whether all pre-treatment coefficients are jointly zero yields $F_{(33,3584)} = 1.32$; $p > F = 0.102$

Table A5. Switching Back to Regular Booking

Dependent variable	Full Sample		Female Sub-Sample		Male Sub-Sample	
	(1)		(2)		(3)	
	Logit - $traditional\_bookable_{it}$		Logit - $traditional\_bookable_{it}$		Logit - $traditional\_bookable_{it}$	
$occupancy\_change\_percent_{it-1}$	0.006	(0.008)	0.012	(0.013)	0.011	(0.008)
$price\_change\_percent_{it-1}$	0.046	(0.081)	0.181	(0.141)	-0.088	(0.128)
$rating\_change\_percent_{it-1}$	-0.259***	(0.051)	-0.141*	(0.065)	-0.534***	(0.108)
$review\_count_{it-1}$	-0.017***	(0.002)	-0.015***	(0.002)	-0.018***	(0.003)
$host\_tenure_{it}$	0.001***	(0.547×10^{-4})	0.001***	(0.782×10^{-4})	0.001***	(0.827×10^{-4})
$host\_is\_superhost_{it}$	0.404***	(0.075)	0.522***	(0.110)	0.228*	(0.115)
$host\_is\_professional_{it}$	-0.136*	(0.060)	0.099	(0.085)	-0.189*	(0.091)
$male_i$	-0.033	(0.049)				
$uncertain_i$	-0.430***	(0.102)				
$gender_i=female_i$ (baseline)						
$White_i$	0.152	(0.084)				
$African American_i$	-0.225*	(0.098)	-0.375***	(0.094)	-0.440***	(0.104)
$Asian_i$	-0.065	(0.109)	0.257*	(0.117)	-0.681***	(0.129)
$Hispanic_i$	0.110	(0.106)	0.213	(0.111)	-0.384***	(0.116)
$race/ethnicity_i=uncertain_i$			0.121	(0.130)	-0.136	(0.124)
$race/ethnicity_i=White_i$						
min_stay_{it}	0.048***	(0.010)	0.038*	(0.016)	0.046***	(0.014)
$photos_{it}$	0.128×10^{-3}	(0.002)	0.005	(0.003)	-0.001	(0.003)
$popularity_{it}$	0.482×10^{-3}	(0.001)	0.001	(0.001)	0.002	(0.001)
$competitor\_instant\_bookable_{it}$	-1.023***	(0.168)	-1.248***	(0.253)	-1.013***	(0.245)
$competitor\_superhost_{it}$	0.141	(0.229)	-0.183	(0.345)	0.572	(0.330)
$competitor\_review\_count_{it}$	0.003	(0.002)	0.004	(0.002)	0.004	(0.003)
$Constant$	-1.715***	(0.138)	-1.251***	(0.190)	-1.981***	(0.207)
Pseudo R-Squared	0.055		0.047		0.068	
Cancellation Policy Dummies	YES		YES		YES	
Room Type Control	YES		YES		YES	
Month Dummies	YES		YES		YES	
# of Obs.	15953		6989		7034	

*** p<0.001, ** p<0.01, * p<0.05; Robust standard errors in parentheses.

Note: I use White hosts as the baseline group in each sub-sample analysis

126

Table A6. App Categories

Category Group	Group Description	Categories in the Group
Content Consumption	Apps that allow consumption of content on a mobile platform	Comics Entertainment Media and Video Music and Audio Sports Video Players
Learn and Explore	Apps that allow users to learn and explore content on the mobile platform	Art and Design Books and References Education Medical News and Magazines Travel and Local Weather Parenting Libraries and Demo
Personal	Apps that are personal to users	Beauty Dating Health and Fitness Lifestyle Personalization Social
Utility – Mobile Access	These are the utility apps that have a major offline/web presence and provides online access through the platform	Business Auto and Vehicles Events Finance House and Home Shopping Transportation
Utility – Mobile Specific	These are utility apps that are present mainly on the mobile platform. Its existence depends on features of smart phone	Maps and Navigation Food and Drink Communication Tools Photography Productivity
Games	These are apps that are categorized as games by Google Play Store	Action Game Adventure Game Arcade Game Board Game Casino Game Casual Game Educational Game Music Game Puzzle Game Racing Game Role Playing Game Simulation Game Sports Game Strategy Game Trivia Game Word Game

Table A7. Summary Statistics

Variables	# Obs.	Mean	Std. Dev.	Min	Max
Key Variables					
$dangerous\_permissions_{it}$	278,955	3.46	2.65	0	22
$essential\_dangerous\_permissions_{it}$	278,955	0.74	0.77	0	10
$nonessential\_dangerous\_permissions_{it}$	278,955	2.71	2.36	0	20
$normal\_permissions_{it}$	278,955	8.21	4.72	0	65
$upgrade_{it}$	278,955	0.48	0.50	0	1
$rating_{it}$	278,955	4.04	0.37	1.86	4.96
$popularity_{it}$	278,955	9.07	15.40	0.00	548.86
App Characteristics					
$rating\_count_{it}$	278,955	150079.30	1108175.00	3	975718104
$dayssinceupdate_{it}$	278,955	176.84	188.52	0	1087
$screenshots_{it}$	278,955	12.50	5.55	2	33
$filesize_{it}$	278,955	24.64	27.89	0.03	1800
$developer\_appcount_{it}$	278,955	29.16	86.82	1	1914
$inAppAdvertising_{it}$	278,955	0.19	0.39	0	1
$inAppPurchases_{it}$	278,955	0.34	0.47	0	1

Note: These statistics are based on the sample before propensity score matching.

Since the computed PageRank scores are numerically small, $popularity_{it}$ is scaled by a factor of 1 million.

	Variables	# Obs.	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
(1)	dangerous_permissions$_{it}$	278,955	1.00													
(2)	essential_dangerous_permissions$_{it}$	278,955	0.50	1.00												
(3)	nonessential_dangerous_permissions$_{it}$	278,955	0.96	0.24	1.00											
(4)	normal_permissions$_{it}$	278,955	0.76	0.37	0.73	1.00										
(5)	upgrade$_{it}$	278,955	-0.04	-0.04	-0.04	0.01	1.00									
(6)	rating$_{it}$	278,955	0.10	0.08	0.09	0.19	0.11	1.00								
(7)	popularity$_{it}$	278,955	0.08	0.04	0.07	0.12	0.11	0.25	1.00							
(8)	rating_count$_{it}$	278,955	0.12	0.05	0.12	0.18	0.03	0.11	0.26	1.00						
(9)	dayssinceupdate$_{it}$	278,955	-0.19	-0.18	-0.15	-0.22	-0.23	-0.17	-0.19	-0.07	1.00					
(10)	screenshots$_{it}$	278,955	-0.11	-0.07	-0.11	-0.08	0.05	0.05	0.07	0.02	0.01	1.00				
(11)	filesize$_{it}$	278,955	0.00	0.03	-0.01	0.02	0.06	0.04	0.06	0.07	0.03	0.30	1.00			
(12)	developer_appcount$_{it}$	278,955	-0.09	-0.02	-0.09	-0.03	-0.03	-0.03	-0.05	-0.02	0.07	0.01	0.01	1.00		
(13)	inAppAdvertising$_{it}$	278,955	-0.02	-0.05	0.00	0.02	0.07	0.05	0.08	0.04	0.02	0.02	0.02	0.00	1.00	
(14)	inAppPurchases$_{it}$	278,955	0.07	0.03	0.07	0.12	0.07	0.20	0.15	0.07	-0.07	0.25	0.29	-0.04	0.13	1.00

Table A9. Moderating Effect of Revenue Models on the Effect of Dangerous Permissions in Delaying Upgrade to Android 6.0

Dependent variable - $upgrade_{it}$	(1)	
	Hazard Ratios – Revenue Model	
$dangerous\_permissions\_ratio_{it}$	0.996***	(0.001)
$inAppAdvertising_{it}$	1.277***	(0.089)
$dangerous\_permissions\_ratio_{it}$ X $inAppAdvertising_{it}$	0.994*	(0.002)
$inAppPurchases_{it}$	1.116	(0.069)
$dangerous\_permissions\_ratio_{it}$ X $inAppPurchases_{it}$	0.999	(0.002)
$\Delta rating\_count_{it}$	1.000***	(0.000)
$rating_{it}$	1.049	(0.032)
$games_i$	0.860***	(0.032)
$personal\_apps_i$	0.926	(0.040)
$utility\_mobilespecific_i$	0.868***	(0.036)
$utility\_mobileaccess_i$	1.059	(0.052)
$learn\_explore_i$	1.090*	(0.047)
$category_i$=$content\_consumption_i$ (baseline)		
$below\_1million_i$	0.916***	(0.022)
$5million\_10million_i$	1.037	(0.039)
$10million\_50million_i$	1.190***	(0.051)
$above\_50million_i$	1.069	(0.103)
$download_i$=$1million\_5million_i$ (baseline)		
$dayssinceupdate_{it}$	0.996***	(0.000)
$screenshots_{it}$	1.012***	(0.002)
$filesize_{it}$	1.002***	(0.000)
Developer Controls	YES	
Log Likelihood	-71414.88	
AIC	142869.8	
BIC	143068.4	
# of Obs.	152,190	
X^2	1486.20	
p	0.000	

*** p<0.001, ** p<0.01, * p<0.05; Clustered standard errors in parentheses.

Dependent Variable	(1) essential_ dangerous_permissions$_{it}$		(2) nonessential_ dangerous_permissions$_{it}$		(3) log(rating$_{it}$)		(4) popularity$_{it}$	
dangerous_permissions_ratio$_{it}$					-0.002	(0.001)	-0.002	(0.013)
Δrating_count$_{it}$	-0.000	(0.000)	-0.000	(0.000)	0.000*	(0.000)	0.000***	(0.000)
rating$_{it}$	0.093	(0.097)	-0.714*	(0.347)			-1.199	(2.437)
upgrade_group$_i$*time$_{t-7}$	0.001	(0.011)	0.039	(0.020)	0.021	(0.021)	0.001	(0.221)
upgrade_group$_i$*time$_{t-6}$	0.012	(0.009)	-0.011	(0.014)	0.003	(0.013)	-0.122	(0.136)
upgrade_group$_i$*time$_{t-5}$	0.001	(0.008)	0.003	(0.014)	-0.020	(0.014)	-0.070	(0.154)
upgrade_group$_i$*time$_{t-4}$	-0.012	(0.010)	0.034*	(0.016)	-0.013	(0.014)	0.119	(0.249)
upgrade_group$_i$*time$_{t-3}$	0.002	(0.010)	-0.005	(0.020)	0.014	(0.017)	0.130	(0.178)
upgrade_group$_i$*time$_{t-2}$	0.007	(0.011)	0.002	(0.020)	0.013	(0.016)	-0.112	(0.203)
upgrade_group$_i$*time$_{t-1}$	-0.012	(0.011)	-0.019	(0.022)	-0.013	(0.015)	-0.023	(0.153)
upgrade_group$_i$*time$_{t0}$			(Omitted Base Case)					
upgrade_group$_i$*time$_{t+1}$	0.005	(0.008)	-0.057**	(0.020)	-0.038*	(0.016)	-0.049	(0.124)
upgrade_group$_i$*time$_{t+2}$	0.004	(0.008)	-0.040*	(0.016)	0.050*	(0.023)	0.327	(0.168)
dayssinceupdate$_{it}$	0.000***	(0.000)	0.001***	(0.000)	-0.000	(0.000)	0.000	(0.000)
screenshots$_{it}$	0.022*	(0.009)	-0.017	(0.011)	0.001	(0.015)	-0.110	(0.139)
filesize$_{it}$	0.000	(0.000)	-0.000	(0.001)	-0.000	(0.000)	-0.007	(0.004)
developer_appcount$_{it}$	-0.001	(0.000)	0.003***	(0.001)	0.002*	(0.001)	-0.021	(0.016)
Constant	0.422	(0.420)	5.327***	(1.417)	139.301***	(0.280)	8.634	(10.456)
Month Dummies	YES		YES		YES		YES	
App Fixed Effect	YES		YES		YES		YES	
Download Bucket Dummies	YES		YES		YES		YES	
F-test (pre-treat coef. jointly 0)	0.71		1.26		0.95		0.31	
p	0.667		0.267		0.465		0.948	
# of Obs.	53672		53672		53672		53672	
R-Squared	0.082		0.031		0.089		0.80	

*** p<0.001, ** p<0.01, * p<0.05

Heteroskedastic Robust Standard Errors in parenthesis

Table A11. Analysis of Effects of Delaying Upgrading to Latest Version – Sub Sample Analysis

Dependent Variable	(1) PSM DID - log(rating$_{it}$)		(2) PSM DID – log(rating$_{it}$)	
dangerous_permissions_ratio$_{it}$	0.003*	(0.001)	-0.002*	(0.001)
Δrating_count$_{it}$	0.000***	(0.000)	0.000***	(0.000)
post_upgrade$_t$	-0.055*	(0.027)	-0.102***	(0.012)
upgrade_group$_i$*post_upgrade$_t$	0.317***	(0.071)	0.116**	(0.039)
late_upgrade_group$_i$*post_upgrade$_t$	-0.226**	(0.073)	-0.055	(0.040)
dayssinceupdate$_{it}$	0.000*	(0.000)	-0.000**	(0.000)
screenshots$_{it}$	-0.051**	(0.018)	0.021**	(0.008)
filesize$_{it}$	-0.002*	(0.001)	0.001**	(0.000)
developer_appcount$_{it}$	0.002**	(0.001)	0.003***	(0.000)
Constant	140.851***	(0.285)	139.605***	(0.153)
Month Dummies	YES		YES	
App Fixed Effect	YES		YES	
Download Bucket Dummies	YES		YES	
Weighted by Population	NO		NO	
# of Obs.	23,562		82,503	
R-Squared	0.113		0.108	

*** p<0.001, ** p<0.01, * p<0.05

Heteroskedastic Robust Standard Errors in parenthesis

Column 1 provides the sub-sample analysis where the treated apps did not reduce seeking non-essential dangerous permissions on average after upgrade, compared to before upgrade.

Column 2 provides the sub-sample analysis where the treated apps reduced seeking non-essential dangerous permissions on average after upgrade, compared to before upgrade.

Table A12. Analysis of Effects of Delaying Upgrading to Latest Version – Moderating effects of Download Groups

Dependent Variable	(1) DID – $essential\_dang\_permissions_{it}$		(2) DID – $nonessential\_dang\_permissions_{it}$		(3) PSM DID – $\log(rating_{it})$		(4) PSM DID – $popularity_{it}$	
$dangerous\_permissions\_ratio_{it}$					0.003***	(0.001)	0.034***	(0.007)
$\Delta rating\_count_{it}$	0.000	(0.000)	-0.000***	(0.000)	0.000***	(0.000)	-0.000	(0.000)
$rating_{it}$	0.248***	(0.047)	-0.384***	(0.102)			11.841***	(1.275)
$post\_upgrade_t$	0.043***	(0.010)	0.036	(0.019)	-0.012	(0.020)	2.481***	(0.324)
$upgrade\_group_i*post\_upgrade_t$	-0.099***	(0.013)	-0.198***	(0.029)	0.003	(0.028)	-1.398**	(0.436)
$upgrade\_group_i*post\_upgrade_t*mediumDG_i$	0.056**	(0.020)	0.097*	(0.041)	0.213***	(0.047)	0.593	(0.471)
$upgrade\_group_i*post\_upgrade_t*lowDG_i$	0.055	(0.098)	0.288	(0.168)	0.324	(0.176)	-1.260	(0.747)
$upgrade\_group_i*post\_upgrade_t*highDG_i$ (baseline)								
$late\_group_i*post\_upgrade_t*mediumDG_i$	0.085***	(0.015)	-0.074*	(0.031)	-0.102**	(0.039)	0.690**	(0.214)
$late\_group_i*post\_upgrade_t*lowDG_i$	0.034	(0.096)	-0.591***	(0.175)	-0.141	(0.177)	0.579	(0.349)
$late\_group_i*post\_upgrade_t*highDG_i$ (baseline)								
$post\_upgrade_t*mediumDG_i$	-0.029**	(0.011)	-0.039	(0.020)	-0.065**	(0.021)	-3.342***	(0.327)
$post\_upgrade_t*lowDG_i$	-0.115***	(0.024)	0.134***	(0.037)	-0.338***	(0.048)	-1.250*	(0.615)
$post\_upgrade_t*highDG_i$ (baseline)								
$dayssinceupdate_{it}$	-0.000	(0.000)	0.000***	(0.000)	-0.000***	(0.000)	-0.005***	(0.000)
$screenshots_{it}$	0.012***	(0.003)	-0.018*	(0.008)	0.025	(0.014)	-0.038	(0.178)
$filesize_{it}$	-0.000	(0.000)	0.002***	(0.001)	-0.000	(0.000)	0.012***	(0.003)
$developer\_appcount_{it}$	-0.000	(0.000)	0.004***	(0.000)	0.003***	(0.000)	-0.015***	(0.002)
$Constant$	-0.171	(0.198)	4.118***	(0.434)	137.919***	(0.208)	-34.361***	(5.882)
Month Dummies	YES		YES		YES		YES	
App Fixed Effect	YES		YES		YES		YES	
Download Bucket Dummies	YES		YES		YES		YES	
Weighted by Population	YES		YES		YES		YES	
# of Obs.	106065		106065		106065		106065	
R-Squared	0.111		0.027		0.176		0.032	

*** $p<0.001$, ** $p<0.01$, * $p<0.05$

Heteroskedastic Robust Standard Errors in parenthesis

The popularity score of the most popular app is over a million times higher than the least popular app in my dataset. Hence it is imperative that apps with large popularity incur numerically larger penalty from delaying upgrade than less popular app.

Bibliography

Abadie, A. 2005. "Semiparametric Difference-in-Differences Estimators," The Review of Economic Studies (72:1), pp. 1-19.

Accenture. 2016. "Accenture Technology Vision," Accenture.

Acquisti, A., and Grossklags, J. 2003. "Losses, Gains, and Hyperbolic Discounting: An Experimental Approach to Information Security Attitudes and Behavior," *2nd Annual Workshop on Economics and Information Security*, pp. 1-27.

Acquisti, A., and Grossklags, J. 2005a. "Privacy and Rationality in Individual Decision Making," *IEEE Security & Privacy* (3:1), pp. 26-33.

Acquisti, A., and Grossklags, J. 2005b. "Uncertainty, Ambiguity and Privacy," *4th Annual Workshop on Economics and Information Security*.

Aguirre, E., Mahr, D., Grewal, D., de Ruyter, K., and Wetzels, M. 2015. "Unraveling the Personalization Paradox: The Effect of Information Collection and Trust-Building Strategies on Online Advertisement Effectiveness," *Journal of Retailing* (91:1), pp. 34-49.

Airbnb. 2017. "Women Hosts and Airbnb: Building a Global Community," Airbnb.

Akerlof, G. A. 1970. "Market for Lemons - Quality Uncertainty and Market Mechanism," *Quarterly Journal of Economics* (84:3), pp. 488-500.

Altonji, J. G., and Blank, R. M. 1999. "Race and Gender in the Labor Market," *Handbook of Labor Economics* (3), pp. 3143-3259.

Ambekar, A., Ward, C., Mohammed, J., Male, S., and Skiena, S. 2009. "Name-Ethnicity Classification from Open Sources," *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*: ACM, pp. 49-58.

Ameri, M., Rogers, S., Schur, L., and Kruse, D. 2017. "No Room at the Inn? Disability Access in the New Sharing Economy," *Academy of Management Discoveries*.

Android. 2016. "Distribution Dashboard," in: *Android*. Retrieved November 10, 2016, from https://developer.android.com/about/dashboards/

Aral, S., and Walker, D. 2011. "Creating Social Contagion through Viral Product Design: A Randomized Trial of Peer Influence in Networks," *Management Science* (57:9), pp. 1623-1639.

Aral, S., Bakos, Y. and Brynjolfsson, E. 2018. "Information Technology, Repeated Contracts, and the Number of Suppliers," *Management Science* (64:2), pp. 592-612.

Austin, P. C. 2010. "Statistical Criteria for Selecting the Optimal Number of Untreated Subjects Matched to Each Treated Subject When Using Many-to-One Matching on the Propensity Score," *American Journal of Epidemiology* (172:9), pp. 1092-1097.

Autor, D. H. 2003. "Outsourcing at Will: The Contribution of Unjust Dismissal Doctrine to the Growth of Employment Outsourcing," *Journal of Labor Economics* (21:1), pp. 1-42.

Ayres, I., Banaji, M., and Jolls, C. 2015. "Race Effects on eBay," *RAND Journal of Economics* (46:4), pp. 891-917.

Ayres, I., and Siegelman, P. 1995. "Race and Gender Discrimination in Bargaining for a New Car," *American Economic Review*, pp. 304-321.

Bakshy, E., Rosenn, I., Marlow, C., and Adamic, L. 2012. "The Role of Social Networks in Information Diffusion," *Proceedings of the 21st International Conference on World Wide Web*: ACM, pp. 519-528.

Bapna, R., Goes, P., Gupta, A., and Jin, Y. 2004. "User Heterogeneity and Its Impact on Electronic Auction Market Design: An Empirical Exploration," *MIS Quarterly* (28:1), pp. 21-43.

Bar-Isaac, H., and Cuñat, V. 2014. "Long-term Debt and Hidden Borrowing," *The Review of Corporate Finance Studies*, 3(1-2), pp. 87-122.

Barach, M.A. 2015 "Search, Screening, and Information Provision: Personnel Decisions in an Online Labor Market," *Working Paper*.

Barkhuus L, Dey A. (2003) Location-Based Services for Mobile Telephony: a study of users' privacy concerns. *Interact.* 3: 702-212.

Barnett, W. P., Feng, M., and Luo, X. 2012. "Social Identity, Market Memory, and First-Mover Advantage," *Industrial and Corporate Change* (22:3), pp. 585-615.

Bastani, S., Blumkin, T. and Micheletto, L. 2015. "Optimal Wage Redistribution in the Presence of Adverse Selection in the Labor Market," *Journal of Public Economics*, *131*, pp. 41-57.

Berry, S., Levinsohn, J. and Pakes, A., 1995. Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society*, pp.841-890.

Bélanger, F., and Crossler, R. E. 2011. "Privacy in the Digital Age: A Review of Information Privacy Research in Information Systems," *MIS Quarterly* (35:4), pp. 1017-1042.

Black, D., Haviland, A., Sanders, S., and Taylor, L. 2006. "Why Do Minority Men Earn Less? A Study of Wage Differentials among the Highly Educated," *The Review of Economics and Statistics* (88:2), pp. 300-313.

Bobbitt-Zeher, D. 2007. "The Gender Income Gap and the Role of Education," *Sociology of Education* (80:1), pp. 1-22.

Brunk, B. 2002. "Understanding the Privacy Space," *First Monday* (7:10).

Buchmueller, T.C. 1995. "Health Risk and Access to Employer-Provided Health Insurance," *Inquiry*, pp.75-86.

Cachon, G.P., Gallino, S. and Olivares, M., 2019. "Does Adding Inventory Increase Sales? Evidence of a Scarcity Effect in US Automobile Dealerships," *Management Science* (65:4), pp.1469-1485.

Cadwalladr, C., and Graham-Harrison, E. 2018. "Revealed: 50 Million Facebook Profiles Harvested for Cambridge Analytica in Major Data Breach," in: *The Guardian*.

Carow, K., Heron, R., and Saxton, T. 2004. "Do Early Birds Get the Returns? An Empirical Investigation of Early-Mover Advantages in Acquisitions," *Strategic Management Journal* (25:6), pp. 563-585.

Chan, J., and Wang, J. 2017. "Hiring Preferences in Online Labor Markets: Evidence of a Female Hiring Bias," *Management Science.* (64:7), pp.2973-2994.

Chang, Y.B., and Gurbaxani, V., 2013. "An Empirical Analysis of Technical Efficiency: The Role of IT Intensity and Competition," *Information Systems Research* (24:3), pp. 561-578.

Chellappa, R. K., and Sin, R. G. 2005. "Personalization Versus Privacy: An Empirical Examination of the Online Consumer's Dilemma," *Information Technology and Management* (6:2-3), pp. 181-202.

Cimpanu, C. 2018. "Google Restricts Which Android Apps Can Request Call Log and Sms Permissions." *ZDNet* Retrieved January 9, 2019, from https://www.zdnet.com/article/google-restricts-which-android-apps-can-request-call-log-and-sms-permissions/

Cornell, B., and Welch I. 1996. "Culture, Information, and Screening Discrimination," *Journal of Political Economy.* (104:3), pp. 542–571.

Cox, M. 2017. "Inside Airbnb: The Face of Airbnb, New York City," Retrieved from http://insideairbnb.com/face-of-airbnb-nyc/

Cui, R., Li, J., and Zhang, D. 2020. "Reducing Discrimination with Reviews in the Sharing Economy: Evidence from Field Experiments on Airbnb," *Management Science* (66:3), pp. 1071-1094

Cullen, Z., and Farronato, C. 2014. "Outsourcing Tasks Online: Matching Supply and Demand on Peer-to-Peer Internet Platforms," *Working Paper*.

Culnan, M. J., and Armstrong, P. K. 1999. "Information Privacy Concerns, Procedural Fairness and Impersonal Trust: An Empirical Investigation," *Organization Science* (10:1), pp. 104-115.

Dawson, G.S., Watson, R.T., and Boudreau, M.C. 2010. "Information Asymmetry in Information Systems Consulting: Toward a Theory of Relationship Constraints," *Journal of Management Information. Systems* (27:3), pp. 143–178.

Dehejia, R. H., and Wahba, S. 2002. "Propensity Score-Matching Methods for Nonexperimental Causal Studies," *Review of Economics and statistics* (84:1), pp. 151-161.

DeSante, C.D. 2013. "Working Twice as Hard to Get Half as Far: Race, Work ethic, and America's Deserving Poor," *American Journal of Political Science.* 57(2), pp. 342–356.

Dewally, M., and Ederington, L. 2006. "Reputation, Certification, Warranties, and Information as Remedies for Seller-Buyer Information Asymmetries: Lessons from the Online Comic Book Market*," The Journal of Business* (79:2), pp. 693-729.

Dewan, S., and Hsu, V. 2004. "Adverse Selection in Electronic Markets: Evidence from Online Stamp Auctions," *Journal of Industrial Economics* (52:4), pp. 497-516.

Dinev, T., and Hart, P. 2006. "An Extended Privacy Calculus Model for E-Commerce Transactions," *Information Systems Research* (17:1), pp. 61-80.

Doleac, J.L., and Stein, L.C. 2013. "The Visible Hand: Race and Online Market Outcomes," *The Economic Journal* (123:572), pp. F469–F492.

Edelman, B. G., and Luca, M. 2014. "Digital Discrimination: The Case of Airbnb.com," *SSRN Working Paper*. http://dx.doi.org/10.2139/ssrn.2377353

Edelman. B., Luca. M., and Svirsky. D. 2017. "Racial Discrimination in the Sharing Economy: Evidence from a Field Experiment," *American Economic Journal: Applied Economics* 9(2), pp. 1–22.

Einav, L., Farronato, C., and Levin, J. 2016. "Peer-to-Peer Markets," *Annual Review of Economics* (8), pp. 615–635.

Eisenmann, T. R., Parker, G., and Van Alstyne, M. W. 2006. "Strategies for Two Sided Markets," *Harvard Business Review,* (84:10), pp. 92.

Eisenmann, T., Parker, G., and Van Alstyne, M. 2011. "Platform Envelopment," *Strategic Management Journal* (32:12), pp. 1270-1285.

Elsen, M., Pieters, R., and Wedel, M. 2011. "When Creativity Meets Repetition: Frequency Effects Depend on Exposure Duration," *ACR North American Advances*.

Fan, Y. 2013. "Ownership Consolidation and Product Characteristics: A Study of the US Daily Newspaper Market," *American Economic Review* (103:5), pp.1598-1628.

Fang, V. W., Tian, X., and Tice, S. 2014. "Does Stock Liquidity Enhance or Impede Firm Innovation?," *The Journal of Finance* (69:5), pp. 2085-2125.

Felt, A. P., Chin, E., Hanna, S., Song, D., and Wagner, D. 2011. "Android Permissions Demystified," *Proceedings of the 18th ACM Conference on Computer and Communications Security*: ACM, pp. 627-638.

Fong, N.M., 2017. "How targeting affects customer search: A field experiment," *Management Science*, (63:7), pp.2353-2364.

Forbes. 2020. "The World's Largest Public Companies." from https://www.forbes.com/global2000/

Fradkin, A. 2015. "Search Frictions and the Design of Online Marketplaces," *Working Paper*.

Fuentelsaz, L., Maicas, J.P. and Polo, Y., 2012. "Switching Costs, Network Effects, and Competition in the European Mobile Telecommunications Industry," *Information Systems Research* (23:1), pp. 93-108.

Gao, H., and Zhang, W. 2016. "Employment Nondiscrimination Acts and Corporate Innovation," *Management Science* (63:9), pp. 2982-2999.

Garber, J. 2018. "GDPR–Compliance Nightmare or Business Opportunity?," *Computer Fraud & Security* (2018:6), pp. 14-15.

Garg R., Telang R. 2013. "Inferring app demand from publicly available data," *MIS Quarterly* (37:4): 1253-1264.

Ghose, A., and Han, S. P. 2014. "Estimating Demand for Mobile Applications in the New Economy," *Management Science* (60:6), pp. 1470-1488.

Goic, M., Guajardo, J. and Ma, L., 2019. "How Does the Value of Geolocation Information Vary Across the Purchase Funnel?," *SSRN Working Paper*.

Goode L. 2018. "App Permissions Don't Tell us Nearly Enough About Our Apps," *Wired* (August 14), https://www.wired.com/story/app-permissions/

Gopal, A., and Gao, G. 2009. "Certification in the Indian Offshore It Services Industry," *Manufacturing & Service Operations Management* (11:3), pp. 471-492.

Grace, M. C., Zhou, W., Jiang, X., and Sadeghi, A.-R. 2012. "Unsafe Exposure Analysis of Mobile in-App Advertisements," *Proceedings of the fifth ACM conference on Security and Privacy in Wireless and Mobile Networks*: ACM, pp. 101-112.

Gradwohl, R. 2018. "Firms Quick to Adopt EU Data Regs Will Have First-Mover Advantage," in: *The Hill*.

Gunter, U., and Önder, I. 2017. "Determinants of Airbnb demand in Vienna and their implications for the traditional accommodation industry," *Tourism Economics.* 1354816617731196.

Gutter, M.S., and Hatcher, C.B. 2008. "Racial Differences in the Demand for Life Insurance," *Journal of Risk and Insurance* 75(3), pp. 677–689.

Halek, M., and Eisenhauer, J. G. 2001. "Demography of Risk Aversion," *Journal of Risk and Insurance* (68:1), pp. 1–24.

Hann, I.-H., Hui, K.-L., Lee, S.-Y. T., and Png, I. P. 2007. "Overcoming Online Information Privacy Concerns: An Information-Processing Theory Approach," *Journal of Management Information Systems* (24:2), pp. 13-42.

Hausman, J.A., 1996. Valuation of new goods under perfect and imperfect competition. In *The economics of new goods* (pp. 207-248). University of Chicago Press.

Hildenbrand J. (2018) What happens when I decline an app permission?. *Android Central* (May 30), https://www.androidcentral.com/what-happens-when-i-decline-app-permission

Hinneburg, A., and Keim, D. A. 1999. "Optimal Grid-Clustering: Towards Breaking the Curse of Dimensionality in High-Dimensional Clustering." *Proceedings of the 25th International Conference on Very Large Databases,* pp. 506-517.

Hoff, K. and Stiglitz, J.E. 1990. "Introduction: Imperfect Information and Rural Credit Markets: Puzzles and Policy Perspectives," *The World Bank Economic Review* (4:3), pp.235-250.

Hoffman, D. L., Novak, T. P., and Peralta, M. 1999. "Building Consumer Trust Online," *Communications of the ACM* (42:4), pp. 80-85.

Hong, Y. and Pavlou, P.A., 2017. "On Buyer Selection of Service Providers in Online Outsourcing Platforms for IT Services," *Information Systems Research* (28:3), pp.547-562.

Hoofnagle, C. J., King, J., Li, S., and Turow, J. 2010. "How Different Are Young Adults from Older Adults When It Comes to Information Privacy Attitudes and Policies?," *SSRN Working Paper*.

Horton, J. J. 2017. "The Effects of Algorithmic Labor Market Recommendations: Evidence from a Field Experiment," *Journal of Labor Economics* (35:2), pp. 345-385.

Hui, K.-L., Teo, H. H., and Lee, S.-Y. T. 2007. "The Value of Privacy Assurance: An Exploratory Field Experiment," *MIS Quarterly* (31:1), pp. 19-33.

Iyer, R., Khwaja, A. I., Luttmer, E. F., and Shue, K. 2009. "Screening in New Credit Markets: Can Individual Lenders Infer Borrower Creditworthiness in Peer-to-Peer Lending?" *SSRN Working Paper*. https://ssrn.com/abstract=1570115

Jensen, C., and Potts, C. 2004. "Privacy Policies as Decision-Making Tools: An Evaluation of Online Privacy Notices," *Privacy Policies as Decision-Making Tools: An Evaluation of Online Privacy Notices*: ACM, pp. 471-478.

Jin, G.Z., and Kato, A. 2006. "Price, Quality, and Reputation: Evidence from an Online Field Experiment," *RAND Journal of Economics* (37:4), pp. 983-1004

Jung E-Y, Baek C, Lee J-D (2012) Product survival analysis for the App Store. *Marketing Letters*. 23(4): 929-941.

Kakar, V., Franco, J., Voelz, J., and Wu, J. 2016. "Effects of Host Race Information on Airbnb Listing Prices in San Francisco," *Journal of Housing Economics* In Press

Kalyanaram, G., and Urban, G. L. 1992. "Dynamic Effects of the Order of Entry on Market Share, Trial Penetration, and Repeat Purchases for Frequently Purchased Consumer Goods," *Marketing Science* (11:3), pp. 235-250.

Kane G., C., and Ransbotham S. 2016. "Research note—Content and Collaboration: An Affiliation Network Approach to Information Quality in Online Peer Production Communities," *Information Systems Research* 27(2), 424-439.

Kanoria, Y., and Saban, D. 2017. "Facilitating the Search for Partners on Matching Platforms: Restricting Agent Actions," *Working Paper*.

Karlsson, L., Kemperman, A., and Dolnicar, S. 2017. "May I Sleep in Your Bed? Getting Permission to Book," *Annals of Tourism Research* (62), pp. 1-12.

Kelley PG, Consolvo S, Cranor LF, Jung J, Sadeh N, Wetherall D. (2012) A Conundrum of Permissions: Installing Applications on an Android Smartphone. *Internat. Conf. Financial Cryptography and Data Security*. 68-79

Kerin, R. A., Varadarajan, P. R., and Peterson, R. A. 1992. "First-Mover Advantage: A Synthesis, Conceptual Framework, and Research Propositions," *The Journal of Marketing* (56:4), pp. 33-52.

King, A. A., Lenox, M. J., and Terlaak, A. 2005. "The Strategic Use of Decentralized Institutions: Exploring Certification with Iso 14001 Management Standard," *Academy of Management Journal* (48:6), pp. 1091-1106.

Kingsley, S. C., Gray, M. L., and Suri, S. 2015. "Accounting for Market Frictions and Power Asymmetries in Online Labor Markets," *Policy & Internet* (4:7), pp. 383-400.

Klein, T. J., Lambertz, C., Spagnolo, G., and Stahl, K. O. 2009. "The Actual Structure of eBay's Feedback Mechanism and Early Evidence on the Effects of Recent Changes," *International Journal of Electronic Business* (7:3), pp. 301-320.

Koopman, C., Mitchell, M., and Thierer, A. 2014. "The Sharing Economy and Consumer Protection Regulation: The Case for Policy Change," *Journal of Business Entrepreneurship and the Law* (8), p. 529.

Kricheli-Katz, T., and Regev, T. 2016. "How Many Cents on the Dollar? Women and Men in Product Markets," *Science Advances*. (2:2), p.e1500599.

Krishnan M. S., Kriebel C. H., Kekre S., Mukhopadhyay T. 2000. "An empirical analysis of productivity and quality in software products," *Management Science*. (46:6): 745-759.

Lee, L., and Ariely, D. 2006. "Shopping Goals, Goal Concreteness, and Conditional Promotions," *Journal of Consumer Research* (33:1), pp. 60-70.

Lehnert, K., Till, B. D., and Carlson, B. D. 2013. "Advertising Creativity and Repetition: Recall, Wearout and Wearin Effects," *International Journal of Advertising* (32:2), pp. 211-231.

Levin, S. 2017. " Airbnb gives in to regulator's demand to test for racial discrimination by hosts," in: *The Guardian*.

Li, Z., Hong, Y., and Zhang, Z. 2016. "Do Ride-Sharing Services Affect Traffic Congestion? An Empirical Study of Uber Entry," *SSRN Working Paper*, http://dx.doi.org/10.2139/ssrn.2838043

Lieberman, M. B., and Montgomery, D. B. 1988. "First‑Mover Advantages," *Strategic Management Journal* (9:S1), pp. 41-58.

Limer E. 2018. "Hundreds of Apps Can Eavesdrop Through Phone Microphones to Target Ads," *Popular Mechanics.* (January 2), https://www.popularmechanics.com/technology/security/a14553262/alphonso-audio-ad-targeting/.

Lin, M., Prabhala, N. R., and Viswanathan, S. 2013. "Judging Borrowers by the Company They Keep: Friendship Networks and Information Asymmetry in Online Peer-to-Peer Lending," *Management Science* (59:1), pp. 17-35.

List, J.A. 2004. "The Nature and Extent of Discrimination in the Marketplace: Evidence from the Field," *Quarterly Journal of Economics* (119:1), pp. 49–89.

Liu, X., and Lynch, L. 2011. "Do Agricultural Land Preservation Programs Reduce Farmland Loss? Evidence from a Propensity Score Matching Estimator," *Land Economics* (87:2), pp. 183-201.

Luca, M. 2017. "Designing Online Marketplaces: Trust and Reputation Mechanisms," *Innovation Policy and the Economy* (17:1), pp. 77-93.

Luo, X., Lu, X. and Li, J., 2019. "When and How to Leverage E-commerce Cart Targeting: The Relative and Moderated Effects of Scarcity and Price Incentives with a Two-Stage Field Experiment and Causal Forest Optimization," *Information Systems Research*, (30:4), pp.1203-1227.

Mahajan, V., and Muller, E. 1998. "When Is It Worthwhile Targeting the Majority Instead of the Innovators in a New Product Launch?," *Journal of Marketing Research*), pp. 488-495.

Maheshwari, S. 2017. "That Game on Your Phone May Be Tracking What You're Watching on TV," in: The New York Times.

Maheswaran, D. and Meyers-Levy, J., 1990. "The influence of message framing and issue involvement," *Journal of Marketing Research*, (27:3), pp.361-367.

Malhotra, N. K., Kim, S. S., and Agarwal, J. 2004. "Internet Users' Information Privacy Concerns (IUIPC): The Construct, the Scale, and a Causal Model," *Information Systems Research* (15:4), pp. 336-355.

Mays, V.M., Cochran, S.D., and Barnes, N.W. 2007. "Race, Race-based Discrimination, and Health Outcomes among African Americans," *Annual Review of Psychology* (58:1), pp. 201–225.

Meyer, B. D. 1995. "Natural and Quasi-Experiments in Economics," *Journal of Business & Economic Statistics* (13:2), pp. 151-161.

Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. 2013. "Distributed Representations of Words and Phrases and Their Compositionality," *Advances in Neural Information Processing Systems*, pp. 3111-3119.

Mitchell, W. 1991. "Dual Clocks: Entry Order Influences on Incumbent and Newcomer Market Share and Survival When Specialized Assets Retain Their Value," *Strategic Management Journal* (12:2), pp. 85-100.

Mnih, A., and Hinton, G. E. 2009. "A Scalable Hierarchical Distributed Language Model," *Advances in Neural Information Processing Systems*, pp. 1081-1088.

Möhlmann, M. 2015. "Collaborative Consumption: Determinants of Satisfaction and the Likelihood of Using a Sharing Economy Option Again," *Journal of Consumer Behaviour* (14:3), pp. 193-207.

Nielsen. 2018. "The Nielsen Total Audience Report : Q3 2018," Nielsen.

Nowak, G. J., and Phelps, J. 1995. "Direct Marketing and the Use of Individual‐Level Consumer Information: Determining How and When "Privacy" Matters," *Journal of Direct Marketing* (9:3), pp. 46-60.

Oestreicher-Singer G, Sundararajan A (2012) Reccomendation Networks and the Long Tail of Electronic Commerce. *MIS Quart*erly (36:1):65-83.

Page L, Brin S, Motwani R, Winograd T (1999) The PageRank citation ranking: Bringing order to the web. *Stanford InfoLab*.

Pager, D., Bonikowski, B., Western, B. 2009. "Discrimination in a Low-age Labor Market: A Field Experiment," *American Sociological Review* (74:5), pp. 777–799.

Pallais, A. 2014. "Inefficient Hiring in Entry-Level Labor Markets," *The American Economic Review* (104:11), pp. 3565-3599.

Park, J., Kim, J., Pang, M.-S., and Lee, B. 2017. "Offender or Guardian? An Empirical Analysis of Ride-Sharing and Sexual Assault," *SSRN Working Paper*. https://ssrn.com/abstract=2951138

Patchin, J. W., and Hinduja, S. 2010. "Trends in Online Social Networking: Adolescent Use of Myspace over Time," *New Media & Society* (12:2), pp. 197-216.

Pavlou, P.A., Liang, H., and Xue, Y. 2007. "Understanding and Mitigating Uncertainty in Online Exchange Relationships: A Principal-Agent Perspective," *MIS Quarterly* (31:1), pp. 105-136.

Pelikan M, Hogben G, Erlingsson U. 2017. "Identifying Intrusive Mobile Apps Using Peer Group Analysis," *Google Security Blog* (July 12), https://security.googleblog.com/2017/07/identifying-intrusive-mobile-apps-using.html

Phelps, J., Nowak, G., and Ferrell, E. 2000. "Privacy Concerns and Consumer Willingness to Provide Personal Information," *Journal of Public Policy & Marketing* (19:1), pp. 27-41.

Pingitore, G., Rao, V., Cavallaro, K., and Dwivedi, K. 2017. "To Share or Not to Share," Deloitte University Press.

Pool, V. K., Stoffman, N., and Yonker, S. E. 2015. "The People in Your Neighborhood: Social Interactions and Mutual Fund Portfolios," *Journal of Finance* (70:6), pp. 2679-2732.

Pope, D.G., and Sydnor, J.R. 2011. "What's in a Picture? Evidence of Discrimination from Prosper.com," *Journal of Human Resources* (46:1), pp. 53-92.

Rainie, L., Kiesler, S., Kang, R., Madden, M., Duggan, M., Brown, S., and Dabbish, L. 2013. "Anonymity, Privacy, and Security Online," *Pew Research Center* (5).

Ramos, J. 2003. "Using Tf-Idf to Determine Word Relevance in Document Queries," *Proceedings of the First Instructional Conference on Machine Learning*, pp. 133-142.

Reimers, C.W. 1983 "Labor Market Discrimination against Hispanic and Black Men," *Review of Economics and Statistics* (65:4), pp. 570–579.

Reynaert, M. and Verboven, F., 2014. "Improving the performance of random coefficients demand models: the role of optimal instruments," *Journal of Econometrics*, (179:1), pp.83-98.

Roberts, M. J., and Schlenkera, W. 2013. "Identifying Supply and Demand Elasticities of Agricultural Commodities: Implications for the US Ethanol Mandate," *The American Economic Review* (103:6), pp. 2265-2295.

Rochet, J.-C., and Tirole, J. 2003. "Platform Competition in Two-Sided Markets," *Journal of the European Economic Association* (1:4), pp. 990-1029.

Romanyuk, G. 2016. "Ignorance Is Strength: Improving the Performance of Matching Markets by Limiting Information." *Working paper*.

Rose, E. 2005. "Data Users Versus Data Subjects: Are Consumers Willing to Pay for Property Rights to Personal Information?," *Proceedings of the 38th Annual Hawaii International Conference on System Sciences*: IEEE, pp. 180c-180c.

Ryan, A. M., Burgess, J. F., and Dimick, J. B. 2015. "Why We Should Not Be Indifferent to Specification Choices for Difference-in-Differences," *Health Services Research* (50:4), pp. 1211-1235.

Sarma, B. P., Li, N., Gates, C., Potharaju, R., Nita-Rotaru, C., and Molloy, I. 2012. "Android Permissions: A Perspective Combining Risks and Benefits," *Proceedings of the 17th ACM Symposium on Access Control Models and Technologies*: ACM, pp. 13-22.

Satariano, A. 2019. " Google Is Fined $57 Million Under Europe's Data Privacy Law," in: *The New York Times*.

Smith, J. A., and Todd, P. E. 2005. "Does Matching Overcome Lalonde's Critique of Nonexperimental Estimators?," *Journal of Econometrics* (125:1-2), pp. 305-353.

Song, T., Yi, C. and Huang, J., 2017. "Whose recommendations do you follow? An investigation of tie strength, shopping stage, and deal scarcity," *Information & Management*, (54:8), pp.1072-1083.

Spiekermann, S., Grossklags, J., and Berendt, B. 2001. "E-Privacy in 2nd Generation E-Commerce: Privacy Preferences Versus Actual Behavior," *Proceedings of the 3rd ACM Conference on Electronic Commerce*: ACM, pp. 38-47.

Stamm, S. M., Tripp; Kuronen, Jessica. 2018. "How Pizza Night Can Cost More in Data Than Dollars," in: *The Wall Street Journal*.

Statista. 2018a. "Android Version Market Share Distribution among Smartphone Owners," in: *Statista*.

Statista. 2018b. "Global Mobile OS Market Share in Sales to End Users from 1st Quarter 2009 to 2nd Quarter 2018," in: *Statista*.

Statista. 2018c. "Smart TV Market Share of Overall TV Market Worldwide from 2015 to 2018," Statista.

Steffensmeier, D., and Demuth, S. 2006. "Does Gender Modify the Effects of Race–Ethnicity on Criminal Sanctioning? Sentences for Male and Female White, Black, and Hispanic Defendants," *Journal of Quantitative Criminology* (22:3), pp. 241-261.

Stiglitz, J. E. 1975. "The Theory of "Screening," Education, and the Distribution of Income," *American Economic Review* (65:3), pp. 283-300.

Stiglitz, J. E. 1977. "Monopoly, Non-Linear Pricing and Imperfect Information: The Insurance Market," *The Review of Economic Studies* (44:3), pp. 407-430.

Stroube, B. 2017. "Tastes and Beliefs: How Economic Incentives Can Encourage Discriminatory Evaluations", Working Paper. London Business School.

Sun, T., Viswanathan, S., and Zheleva, E. 2019. "Creating Social Contagion through Firm Mediated Message Design: Evidence from a Randomized Field Experiment," *Management Science* Forthcoming.

Tang, C., Ross, K., Saxena, N., and Chen, R. 2011. "What's in a Name: A Study of Names, Gender Inference, and Gender Behavior in Facebook," *Database Systems for Advanced Applications* (6637), pp. 344-356.

Taylor V. F., and Martinovic I. 2016. "Securank: Starving permission-hungry apps using contextual permission analysis," *Proceedings of the 6th Workshop on Security and Privacy in Smartphones and Mobile Devices*, 43-52.

Tsai, J. Y., Egelman, S., Cranor, L., and Acquisti, A. 2011. "The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study," *Information Systems Research* (22:2), pp. 254-268.

Tucker, C.E., 2014. "Social networks, personalized advertising, and privacy controls," *Journal of Marketing Research* (51:5), pp.546-562.

Tversky, A. and Kahneman, D., 1981. "The Framing of Decisions and the Psychology of Choice," *Science*, (211:4481), pp.453-458.

Varian, H. R. 2009. "Economic Aspects of Personal Privacy," in *Internet Policy and Economics*. Springer, pp. 101-109.

Wang, D., Li, M., Guo, P., and Xu, W. 2016. "The Impact of Sharing Economy on the Diversification of Tourism Products: Implications for Tourist Experience," in *Information and Communication Technologies in Tourism 2016*. Springer, pp. 683-694.

Wang, D., Xi, S., and Gilheany, J. 2015. "The Model Minority? Not on Airbnb.com: A Hedonic Pricing Model to Quantify Racial Bias against Asian Americans," *Technological Science* 2015090104.

Watanabe T, Akiyama M, Kanei F, Shioji E, Takata Y, Sun B, Ishi Y, Shibahara T, Yagi T, Mori T. (2017) Understanding the origins of mobile app vulnerabilities: A large-scale measurement study of free and paid apps. *Proc. 14th Internat. Conf. Mining Software Repositories*. 14-24.

Wei, X., Gomez, L., Neamtiu, I., and Faloutsos, M. 2012. "Permission Evolution in the Android Ecosystem," *Proceedings of the 28th Annual Computer Security Applications Conference*: ACM, pp. 31-40.

Weiss, G.N., Pelger, K., and Horsch, A. 2010. "Mitigating Adverse Selection in P2P lending – Empirical Evidence from Prosper.com," *SSRN Working Paper*, https://ssrn.com/abstract=1650774

Wikipedia. 2020. "Platform Economy." Retrieved July 9, 2020, from https://en.wikipedia.org/wiki/Platform_economy

WSJ. 2017. "The Billion Dollar Startup Club." Retrieved December 18, 2017, from http://graphics.wsj.com/billion-dollar-club/

Xue, L., Ray, G. and Sambamurthy, V., 2012. "Efficiency or Innovation: How do Industry Environments Moderate the Effects of Firms' IT Asset Portfolios?" *MIS Quarterly* (36:2), pp. 509-528.

Yang, Z., Li, M., and Ai, H. 2006. "An Experimental Study on Automatic Face Gender Classification," *Pattern Recognition, 2006. ICPR 2006. 18th International Conference on*: IEEE, pp. 1099-1102.

Ye, S., Gao, G., and Viswanathan, S. 2014. "Strategic Behavior in Online Population Systems: Evidence from Revoking on eBay," *MIS Quarterly* (38:4), pp. 1033-1056.

Ye, S., Viswanathan, S., and Hann, I.-H. 2018. "The Value of Reciprocity in Online Barter Markets: An Empirical Investigation," *MIS Quarterly* (42:2), pp. 521-549.

Zervas, G., Proserpio, D., and Byers, J. W. 2014. "The Rise of the Sharing Economy: Estimating the Impact of Airbnb on the Hotel Industry," *Journal of Marketing Research* (54:5), pp. 687–705.

Zimmer, J. C., Arsal, R. E., Al-Marzouq, M., and Grover, V. 2010. "Investigating Online Information Disclosure: Effects of Information Relevance, Trust and Risk," *Information & Management* (47:2), pp. 115-123.