

PEER EFFECTS IN PRODUCT ADOPTION*

Michael Bailey[†] Drew Johnston[‡] Theresa Kuchler[§] Johannes Stroebel[¶] Arlene Wong^{||}

Abstract

We study the nature of peer effects in the market for new cell phones. Our analysis builds on de-identified data from Facebook that combine information on social networks with information on users' cell phone models. To identify peer effects, we use variation in friends' new phone acquisitions resulting from random phone losses and carrier-specific contract terms. A new phone purchase by a friend has a large positive and long-term effect on an individual's own demand for phones of the same brand, most of which is concentrated on the particular model purchased by the friend. We provide evidence that social learning contributes substantially to the observed peer effects. While peer effects increase the overall demand for cell phones, a friend's purchase of a new phone of a particular brand can reduce individuals' own demand for phones from competing brands—in particular those running on a different operating system. We discuss the implications of these findings for the nature of firm competition. We also find that stronger peer effects are exerted by more price-sensitive individuals. This positive correlation suggests that the elasticity of aggregate demand is substantially larger than the elasticity of individual demand. Through this channel, peer effects reduce firms' markups and, in many models, contribute to higher consumer surplus and more efficient resource allocation.

JEL Codes: L1, L2, M3, D4

Keywords: Peer Effects, Demand Spillovers, Social Learning

*This version: May 9, 2019. For their helpful comments, we thank Tim Armstrong, John Campbell, Liran Einav, Matt Gentzkow, Ed Glaeser, Paul Goldsmith-Pinkam, Michal Kolesar, Matthew Notowidigdo, Alex Peysakhovich, Luigi Pistaferri, Paulo Somaini, and Chris Tonetti, as well as seminar and conference participants at Baruch, Carnegie Mellon, Facebook, NYU, Stanford, UNC Chapel Hill, UT Austin, Wharton, Yale, and the Empirical Macro Workshop in Las Vegas. We thank the Center for Global Economy and Business at NYU Stern for generous research support. We also thank Abhinav Gupta, Sung Lee, and Hongbum Lee for excellent research assistance. This research was facilitated through a research consulting agreement between some of the academic authors (Johnston, Kuchler, and Stroebel) and Facebook. Bailey is an employee at Facebook.

[†]Facebook. Email: mcb Bailey@fb.com

[‡]New York University, Stern School of Business. Email: drew.johnston@nyu.edu

[§]New York University, Stern School of Business. Email: tkuchler@stern.nyu.edu

[¶]New York University, Stern School of Business, NBER, and CEPR. Email: johannes.stroebel@nyu.edu

^{||}Princeton University and NBER. Email: arlenewong@princeton.edu

Many consumption decisions are influenced by related decisions made by a person's peers. Such peer effects have important implications for firms and policy-makers. For example, in the presence of peer effects, the aggregate elasticity of demand may be lower than the elasticity of individual demand, which would reduce firms' markups and increase consumer surplus. Peer effects can also increase a firm's new-customer acquisition value beyond the direct effect of this customer on profits. From a macro perspective, peer effects in consumption imply that the effects of stimulus policies on aggregate demand are larger than those estimated from directly-affected individuals. Despite the economic importance of such peer effects, there is limited empirical evidence on their exact nature and the resulting implications. For example, peer effects may lead an individual to buy a new phone when his friend gets a new phone, but the effect of this purchase on firm profits depends on whether it represents incremental demand or the retiming of an already planned purchase. The implications of peer effects also depend on whether any resulting changes in demand are restricted to the precise model purchased by the peer, or whether there are positive or negative demand spillovers to competing brands.

In this paper, we explore the nature of peer effects in the U.S. cell phone market. We find that peer effects are large and heterogeneous, with price-sensitive individuals exerting particularly great influence on their friends. This positive correlation between social influence and price sensitivity implies a substantial wedge between the respective elasticities of individual and aggregate demand, and suggests that peer effects reduce firms' markups. Peer effects are long-lasting and generate substantial incremental demand. Positive peer effects are largest for the model purchased by the peer, and they display some positive demand spillovers to other devices of the same brand, most likely through a social learning channel. The size of peer effect on same-brand demand often exceeds that of the effect on total phone demand, suggesting that some new same-brand purchases come at the expense of purchases from competing brands—in particular those on different operating systems. This finding highlights the important role that peer effects have on the nature of competition between firms.

We work with de-identified data from Facebook, the world's largest online social networking site. Facebook has over 242 million active users in the U.S. and Canada, and more than 2.3 billion users globally. In this data set, we observe individuals' social networks as represented by their Facebook friends, which have been shown to provide a fair representation of real-world U.S. friendship networks. For mobile active users, we also observe data on the device model used to log into their Facebook accounts. These data allow us to identify the timing of new phone acquisitions. We combine these data sets to explore how phone purchases by a user's friends influence the user's own phone-purchasing behavior.

To identify peer effects separately from common shocks or common preferences within friendship groups, we exploit quasi-random variation in friends' phone purchasing behaviors. Useful sources of variation need to shift a friend's probability of acquiring a new phone in a given week, without affecting the probability of a user herself purchasing a new phone through any channel other than peer effects. We use two separate sources of variation that fit these requirements. Firstly, we use the number of friends who break or lose their phones in a given week to instrument for the number of

friends who purchase a new phone in that week. The identifying assumption is that how many friends break or lose their phones in a given week is conditionally random and unrelated to a user's own propensity to buy a new phone in that week. We provide various pieces of evidence in support of this assumption. We identify individuals who randomly break or lose their phones by applying natural language processing and machine learning techniques to the universe of public posts on Facebook. This choice allows us to detect posts, such as *"Phone broken...Ordered a new one but if anyone needs me urgently, call Joe,"* which signal the random phone loss by a peer. We show that people are substantially more likely to obtain a new phone in the week after posting such messages. Our second instrument for the number of friends who obtain a new phone in a given week is the number of peers who are likely eligible for a contract renewal in that week, which is often aligned with an upgrade to a new device.

We improve the power of these instruments by exploiting variation in not only how many friends experience the conditionally random event, but also which friends do so. Specifically, for both instruments, we use neural networks to estimate the probability that each individual would obtain a new phone conditional on the event, exploiting, for example, the fact that older individuals are more likely to buy a new phone immediately after breaking their old device. Our final instrument is the sum of these estimated propensities across all individuals who experience the event, controlling for the distribution of these propensities in the overall pool of friends (e.g., conditional on the average age in a person's friendship network, is it her old or young friends who break their phones this week).

Across both instruments, we obtain peer effect estimates of similar magnitude. Having one additional friend who purchases a new phone in a given week increases an individual's own probability of buying a new phone in the following week by 0.041 and 0.026 percentage points, estimates obtained using the random phone loss instrument and the contract renewal instrument, respectively. These estimated effects are large relative to the weekly probability of buying a new phone of about one percentage point. Interestingly, these instrumental variables (IV) estimates are similar in size to the ordinary least squares (OLS) estimates. One interpretation of this finding is that, in our setting, common shocks and common preferences may not lead to a large bias of the OLS estimate.

In addition to exploring the immediate response of an individual's own purchasing behavior to new phone acquisitions by her friends, we also analyze the extent to which this situation generates new purchases instead of pulling forward already-planned future purchases. We find that a random phone loss by an individual has a positive effect on the total number of phones purchased by her friends in each of the following ten months—though the magnitude of this effect starts to decline after about three months. Peer effects thus cause an increase in the total number of phone purchases, at least over intermediate horizons. One immediate and important implication of this creation of new demand through peer effects is that the value of acquiring new customers exceeds the direct effect of these customers on the revenue and profitability of the firm.

In the next step, we explore heterogeneities in peer effects along characteristics of potential influencers and potentially-influenced individuals. We focus on heterogeneity in the local average treat-

ment effects of the random phone loss instrument, which has the most power in the baseline specification. Closer friends and family friends exert particularly large peer effects. We also find large heterogeneities in the peer effects exerted by different demographic groups, but we find little variation in individual susceptibility to influence along the same demographic characteristics. For example, younger and less-educated individuals have the largest effects on their friends' purchasing behaviors, but individuals in these groups are no more likely to be influenced by phone purchases of their friends. These heterogeneities in peer influence have important implications for the optimal design of "viral" or seed marketing campaigns, which target marketing activity to a small set of early adopters who will generate follow-on demand through peer effects (see, for example, Watts, Peretti, and Frumin, 2007).

Importantly, we also find that those individuals who exert larger peer effects are generally more price sensitive, measured as the effect of a price cut for a phone model on the probability of purchasing that model. This result suggests that the difference between the elasticities of aggregate and individual demand induced by peer effects is even larger than implied by the average peer effect. This higher price elasticity faced by firms leads to lower optimal markups than would otherwise prevail. The positive correlation between price sensitivity and peer influence can also provide an explanation for the sometimes-puzzling observation that many markets with per-period supply constraints clear through queuing rather than through price increases. If price increases disproportionately reduce demand from those individuals who have large peer effects on their friends, then an optimal dynamic pricing strategy might be willing to accept lower revenue today in return for additional sales generated through peer effects in future periods. While demand at lower prices may sometimes exceed supply, as it often does around the release dates of new iPhones, assignment via queuing is likely to disproportionately select particularly enthusiastic individuals who will exert the largest peer effects. Similar mechanisms might be at work in other settings where limited supply is assigned through queuing that can help select individuals who will exert particularly large peer effects and thus generate subsequent sales (e.g., new sneakers, restaurants, Broadway shows, or the famous Cronuts).

The role of peer effects in increasing the elasticity of aggregate demand above the elasticity of individual demand also has important implications for evaluating the effects of mergers between firms. One standard approach in the literature is to use the model-implied equilibrium under a particular market structure jointly with demand estimates from the data to infer marginal costs (e.g., Nevo, 2001). The recovered marginal costs and estimated demand parameters are then used to determine the new equilibria that would result from mergers of firms. These models typically ignore peer effects. A model with peer effects would imply a lower optimal markup. Therefore, an econometrician that ignores peer effects would overestimate marginal costs and the anticompetitive effects of a merger.

In the second part of the paper, we explore whether peer effects are limited to the brand purchased by the peer, or whether there are demand spillovers to other brands. To do so, we first predict the probability that each individual would purchase a phone in each of three broad brand categories (iPhone, Galaxy, other). We then exploit variation in this probability among friends who randomly break their

phones in a given week (conditional on the average of this probability among all friends) to instrument for the number of friends that purchase phones of that particular brand. The identification assumption is similar to before: conditional on the characteristics of all friends and other controls, it is random whether, in a given week, the friends who lose their phones are those who are more likely to replace it with a new iPhone or those who are likely to purchase a new Samsung Galaxy.

There are three main take-aways from the cross-brand analysis. Firstly, for all three brand categories, positive peer effects are largest for phones in the same category as that purchased by the peer. Secondly, these same-brand peer effects are largest for less-well-known but cheaper “other” phones, and they are smallest for the expensive and well-known iPhones. These facts suggest that social learning is an important part of the explanation for these peer effects, since social learning should be more important for lesser-known brands, while “keeping up” effects should be more important for expensive brands that can signal high status. Our findings thus contribute to a literature that both establishes informational frictions as an impediment to the adoption of new products and technologies, and argues that information provisions by peers can overcome these frictions (e.g., Foster and Rosenzweig, 1995; Duflo and Saez, 2003; Mobius, Niehaus, and Rosenblat, 2005; Conley and Udry, 2010).

The third main take-away relates to across-brand demand spillovers. In particular, we find that when a friend purchases a new phone, this event increases a user’s own propensity of purchasing a phone from a competing brand on the same operating system, while reducing their propensity of purchasing a phone from a competing brand on a different operating system. In other words, while some of the observed positive same-brand peer effects arise by generating entirely new purchases, others come from pulling demand away from rival firms with competing operating systems. This finding shows that peer effects have important competitive implications for firms: losing a customer to a competitor does not only mean missing out on positive peer effects that this customer could have had, but will also lead to future losses of other customers through competitive peer effects. These implications of peer effects for the demand of competitors’ brands complement a large literature that has explored similar spillover effects of advertising (e.g., Roehm and Tybout, 2006; Sahni, 2016; Shapiro, 2018; Sinkinson and Starc, 2018). In that literature, researchers regularly find positive demand spillovers to non-advertised competitor brands. Our finding of frequent and substantial negative across-brand peer effects highlights that the implications of peer effects for the competitive dynamics between firms can be qualitatively different to those from the spillover effects of marketing activities.

In the final analysis, we explore the extent to which peer effects are specific to the model purchased by the peer. While we are able to predict with some accuracy which brand an individual will purchase (e.g., iPhone vs. Galaxy), we have less capability to predict the exact phone model (e.g., iPhone 6 vs. iPhone 6s). This disparity means that our previous identification strategy, which exploited small-sample variation in such propensities among people with random phone losses, does not translate to our model-level analysis. Therefore, we use OLS specifications for this part of the analysis, controlling for a large number of observable characteristics of an individual and their peers. While common shocks

and preferences may lead to upward-biased estimates on average, the across-device heterogeneity in estimates that we explore is unlikely to be the result of such biases.

Our model-level analysis reveals a number of interesting patterns. Firstly, the peer effects on the probability of buying a phone of the same model as that purchased by the peer (same-model peer effects) are substantially larger than the effects on the probability of buying a phone of a different model (different-model peer effects). Secondly, the magnitude of same-model peer effects does not vary with the cost of the model purchased by the peer, but it is decreasing in the time since the model release. These findings provide further evidence for a social learning channel behind the observed peer effects, since information acquired from peers is likely of particular relevance to the specific model bought by the peer but would become less important as the model becomes more well-known over time. The third pattern is that same-brand, different-model peer effects are more than twice as large as different-brand peer effects. Among the main device manufacturers, these same-brand, different-model peer effects are largest for Apple, which co-brands all of its devices under the iPhone brand, and they were smallest for LG, which does not do so. This result provides evidence for the value of umbrella branding strategies, which can channel spillovers from peer effects towards other models of the same brand (see Erdem, 1998; Erdem and Sun, 2002).

Our paper contributes to a literature that has studied the role of peer effects in a wide range of economic and financial decisions. Peers have been shown to influence consumption choices (e.g., Mobius, Niehaus, and Rosenblat, 2005; Kuhn et al., 2011; Moretti, 2011; De Giorgi, Frederiksen, and Pistaferri, 2016; Han, Hirshleifer, and Walden, 2016) and a variety of household financial decisions, such as saving for retirement (e.g., Duflo and Saez, 2003; Beshears et al., 2015), investment choices (e.g., Hong, Kubik, and Stein, 2004; Bursztyn et al., 2014; Ouimet and Tate, 2017), housing market decisions (e.g., Bailey et al., 2017, 2018b), and charitable giving (e.g., DellaVigna, List, and Malmendier, 2012). Peer effects also play an important role in explaining education decisions (e.g., Hoxby, 2000; Sacerdote, 2001, 2011) and labor market outcomes (Mas and Moretti, 2009). Prior work has studied peer effects in product and technology adoption decisions; one focus of this literature has been how social learning can help the diffusion of new technologies in developing countries (e.g., Foster and Rosenzweig, 1995, 2010; Conley and Udry, 2010; Oster and Thornton, 2012; Kremer and Miguel, 2007). In the developed world, peer effects have been shown to affect the adoption of new technologies such as solar panels (e.g., Bollinger and Gillingham, 2012; Allcott and Kessler, 2019).¹ Within the literature that has studied peer effects in product adoption decisions, we are the first, to our knowledge, to identify important competitive spillovers to other models and brands. Our setting and research design also allow us to expand our understanding of peer effects along other dimensions. For example, we are able to document that peer effects can generate additional demand rather than just a retiming of demand. We can also identify characteristics of particularly influential individuals, as well as the correlation of peer influence with price sensitivity, which has important implications for firms' price-setting decisions.

¹A vast, related literature has studied the diffusion process of beliefs or information through social networks (e.g., Valente, 1996; Jackson and Yariv, 2005; Young, 2009; Christakis and Fowler, 2013; Bakshy et al., 2012).

1 Data Description

A central challenge for studying peer effects in product adoption decisions is the need to observe both social networks and product adoption behavior within the same data set. We overcome this measurement challenge by exploring peer effects in phone purchasing decisions using de-identified data from Facebook, the world’s largest online social networking site with 242 million monthly active users in the U.S. and Canada. In the U.S., Facebook primarily serves as a platform for real-world friends and acquaintances to interact online, and people usually only add connections to individuals on Facebook whom they know in the real world (Jones et al., 2013). As a result, friendships on Facebook provide a good approximation of real-world friendship networks (see Bailey et al., 2018a).

For each Facebook user, we observe basic demographic information such as their date of birth, gender, and county location, as well as the set of individuals that they are connected to. Using the language adopted by the Facebook community, we call these connections “friends.” The vast majority of Facebook users regularly access their Facebook accounts from their cell phones.² For these mobile active users, we observe data on the cell phone carrier and the phone model used to access the Facebook app. We use these data to identify when a user obtains a new phone.³ Since we can only observe a new phone model when the user logs into the Facebook app for the first time from the new device, we can generally pinpoint the timing of the purchase to roughly the week that a new device is acquired. Our unit of observation is therefore the purchasing behavior of a user in a given week.

In our analysis, we focus on U.S.-based Facebook users between 18 and 65 years of age who have between 100 and 1,000 friends on Facebook. We also require users to access Facebook on their phones across two consecutive weeks in order to be able to observe the timing of potential phone purchases. Our primary sample covers the purchasing behavior of these individuals across four consecutive weeks in May 2016. These weeks were chosen to be relatively far away from both major phone release dates and major shopping holidays (such as Black Friday or Cyber Monday), which could confound our estimates. We are left with a sample of about 335 million user-weeks.

Table 1 provides summary statistics on our sample. The average user in our sample is 35 years old, with a 10th–90th-percentile age range of 21 years to 53 years. Roughly 58% of users in our sample are male. Fifty-five percent of the users have an iPhone and 27% have a Samsung Galaxy; the rest of the users are relatively fragmented across many other phone models. The average user has a phone that is 386 days old, while the median user has a phone that is slightly less than one year old. The 10th–90th-percentile range of phone age is between 63 days and 770 days. About 0.95% of all users acquire a new phone in a given week. The average user has 328 friends in the sample as well as about 3.1 new phone purchases among friends in a given week.

²Facebook reports in its July 26, 2018, 10-Q filing: “Substantially all of our daily and monthly active users [...] access Facebook on mobile devices.”

³The process of determining when a user obtains a new phone involves a number of steps, including the removal of likely work phones or phones borrowed from a friend, as well as dropping temporary phones with only a few log-ins. Because Facebook only records the device model but no unique device identifier, we are unable to detect switches between two devices of the same model.

Table 1: Summary Statistics

	Mean	St. Dev.	P10	P25	P50	P75	P90
User Characteristics							
Age (Years)	35.3	12.1	21	25	33	44	53
Male	0.58	0.49	0	0	1	1	1
Phone Age (Days)	385.7	322.4	63	151	315	542	770
Buys Phone (%)	0.95	9.70	0	0	0	0	0
Has iPhone	0.55	0.50	0	0	1	1	1
Has Galaxy	0.27	0.44	0	0	0	1	1
Friend Characteristics							
Friends in Sample	328.0	205.3	125	168	264	433	642
Friends with Phone Purchases	3.12	2.95	0	1	2	4	7
Friends with Public Statuses	53.9	42.0	17	26	41	67	107
Friends Posting about Breaking/Losing Phone	0.27	0.65	0	0	0	0	1
Friends at Phone Age Threshold	1.86	1.86	0	0	1	3	4

Note: Table presents summary statistics for our baseline panel. The unit of observation is a user-week, and our data consist of approximately 335 million such user-weeks. For each characteristic, we present the mean, standard deviation, and the 10th, 25th, 50th, 75th, and 90th percentiles of the distribution.

2 Research Design

We next outline how we use the data described above to identify peer effects in cell phone-purchasing behavior. Our most basic specification seeks to understand a Facebook user’s decision to buy a new phone in a given week as a function of the prior or contemporaneous purchases of her friends. The challenge for identifying such peer effects is that individuals tend to be friends with others who are similar to them across many dimensions (McPherson, Smith-Lovin, and Cook, 2001; Bailey et al., 2018a,b). For example, in the context of our study, an Apple enthusiast may primarily be friends with other Apple enthusiasts. Even in the absence of peer effects, these friends may thus have similar phone-purchasing patterns, such as buying a new iPhone around its release date. As a result, observing a correlation in purchasing behavior within friendship groups does not necessarily provide evidence for peer effects (see Manski, 1993, for an extended discussion).

Our approach to solving this identification challenge is to develop instrumental variables for the purchasing behavior of a person’s friends. A successful instrument should shift the purchasing behavior of a person’s friends without affecting the purchasing behavior of that person through any channel other than peer effects. We propose two instruments that meet this exclusion restriction: first, the number of a user’s friends who randomly lose their phones, and second, the number of friends who have owned their phones for exactly two years, and whose contract is thus likely up for renewal.

2.1 Random Phone Loss Instrument

Our first instrument is based on the idea that individuals are substantially more likely to buy a new phone in a week in which they lose or break their current phone. Provided that when a friend breaks

or loses her current phone, this event only influences the probability that a user herself purchases a new phone through peer effects, the number of friends who experience such shocks can then be used to instrument for the number of friends who purchase new phones.

Figure 1: Sample Posts About Randomly-Lost Phones



The first step in constructing this instrument is to determine which individuals randomly break or lose their phones in a given week. We do so by analyzing public posts on Facebook that relate to such events. Figure 1 provides examples of such posts, which were relatively common during our sample period, since users regularly posted on Facebook to explain to their friends why they were not returning calls or text messages. We use a machine learning-based approach to classify the universe of public Facebook posts in a given week, allowing us to assign an indicator $\mathbb{1}(RandomPhoneLoss_{i,t})$ to individuals who post about a random phone loss in that week.⁴ Specifically, we use word embeddings and convolutional neural networks to create a classifier that can process the semantic content of text (see Mikolov, Yih, and Zweig, 2013, and Appendix A.1 for details).⁵ This approach allows us to identify posts such as “*R.I.P phone. You will be missed*” that would be difficult to capture with regular expression searches.⁶ Using the neural network classifier, we identify around 67,000 public posts about broken or

⁴We only have access to posts from individuals who have set their privacy settings for that specific post to “public” at the time of the analysis, rendering the post visible to any individual with the URL. Table 1 shows that while the average person has about 328 friends in total, only about 54 of those members have set their statuses as public.

⁵Word embeddings are commonly used in natural language processing applications and provide a way of expressing the meaning of a word or phrase as a multidimensional (often 100- to 500-dimensional) vector in a way that preserves semantic and syntactic relationships with other words. Word embeddings are trained in an unsupervised manner, using a large corpus of text (in our case, the entirety of Wikipedia) as a training set to uncover relationships between terms. To classify the content of a post, we concatenate vectors representing each word in the post to form a matrix. We then feed the resulting matrix into a convolutional neural network to classify whether the post relates to a user randomly breaking or losing his phone. This neural network classifier is trained on a manually-classified sample of public posts.

⁶We have also implemented a model using a regular expression-based classifier, which produced an instrument that had less power but found largely similar results. This simpler classifier is used to reinforce our main model in an approach

lost phones per week. Table 1 shows that, in a given week, the average person has 0.27 friends who publicly post about losing their phones.

Panel A of Figure 2 visualizes the first stage of the random phone loss instrument. It shows the probability of purchasing a new phone in each week, splitting individuals according to their posting behavior in week 0. The green-triangle line corresponds to individuals who publicly post about a random phone loss in week 0. The orange-circle line corresponds to individuals with a public post that was not about a random phone loss, and the blue-square line corresponds to individuals without a public post in week 0. In the weeks prior to posting about a random phone loss, the purchasing behavior of individuals who post about such a phone loss in week 0 has a broadly similar trend to that of other individuals, although it has a somewhat higher level. (As we describe below, our research design will account for this higher level). In week 0, those individuals who posted about a random phone loss have a substantially higher probability of acquiring a new phone. Specifically, in the week they posted about losing their phone, about 10% of individuals with a post identified by our classifier get a new phone. The probability of purchasing a new phone remains slightly elevated in the week following the post about the random phone loss before returning to its baseline rate.

While the probability of getting a new phone spikes in the week of the post and remains elevated in the following week, the sum of these probabilities is far below 100%, meaning that we do not observe a new phone purchase for every individual whom we identify as posting about a random phone loss. There are several reasons for this result. Firstly, our classifier is likely to include some “false positive” posts that we incorrectly identify as indicating a random phone loss. For example, our classifier cannot perfectly separate posts that mention that someone’s “phone is dead” into those that talk about a dead battery and those that talk about a permanently broken phone.⁷ A second explanation is that some users may continue to use a phone with a broken screen or damage of another type. Users may also be able to repair broken phones or recover lost or stolen phones. Finally, our data do not allow us to identify individuals who replace a broken phone with a new phone of the exact same model. In these instances, however, peer effects are likely to be small, and not observing these switches is unlikely to substantially bias our results.

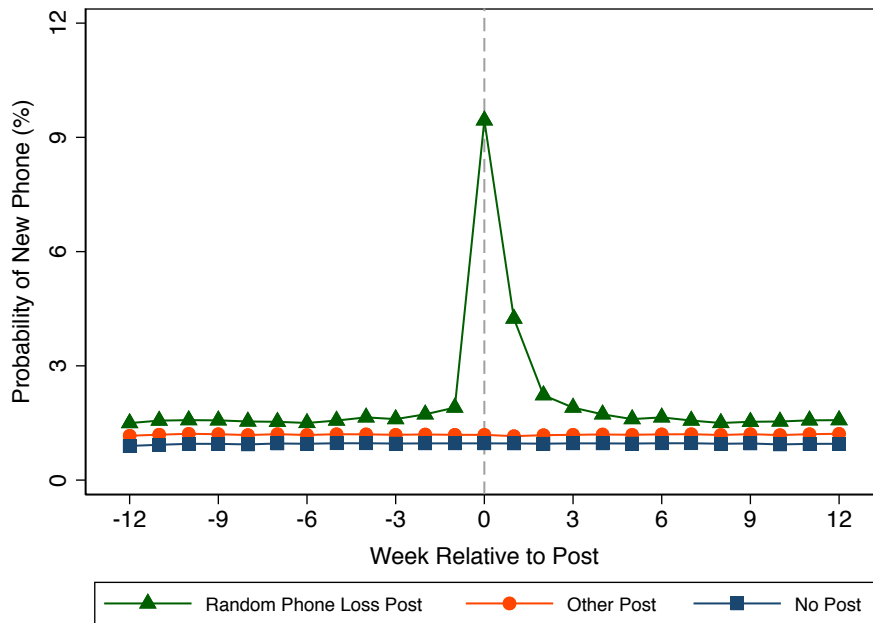
Based on this classification of a random phone loss, a basic identification strategy would instrument for the number of friends who purchase a new phone in a given week with the number of friends who publicly post about randomly breaking or losing their phones in that week. The associated identifying assumption would be that the number of friends losing or breaking their phones in a given week is conditionally random. To strengthen the validity of this exclusion restriction, we include a number of controls in specifications using this first instrument. One possible concern is that the purchasing be-

inspired by ensemble classifiers. Further information about the regex-based classifier can be found in Section A.1.

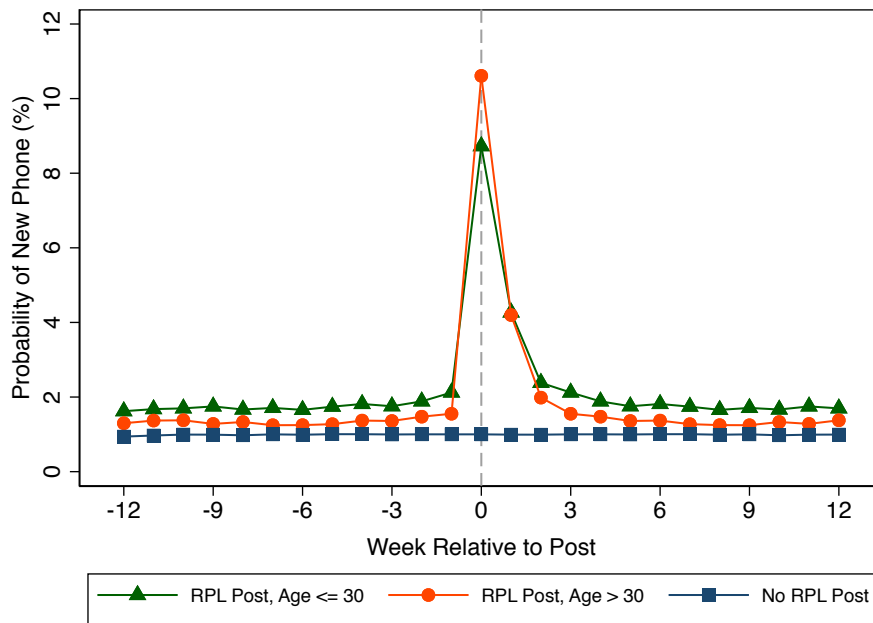
⁷Properly weighting “false positives” and “false negatives” was an important consideration when constructing our classifier, and we chose a threshold that balanced the number of the posts found with the conditional probability of switching of the posters. We also trained an alternative classifier that was better at rejecting false positives and gave a conditional $Pr(BuysPhone_{i,t} | \mathbb{1}(RandomPhoneLoss_{i,t}))$ of 13.4%, although the number of posts found decreased by 85%. This decrease in the number of false positives thus weakened our instrument.

Figure 2: Random Phone Loss Instrument

(A) Probability of New Phone, by Week



(B) Probability of New Phone, by Week (Split by Age)



Note: Panel A shows the probability of purchasing a new phone in every week, splitting users by their posting behavior in week 0. The line *Random Phone Loss Post* shows the behavior of users who have a public post in week 0 that relates to a random phone loss. The line *Other Post* captures the behavior of those who have a public post in week 0 that does not relate to a random phone loss, while the line for *No Post* tracks the behavior of those individuals without a public post in week 0. Panel B shows the probability that a user of each age group buys a phone in the weeks after posting about randomly losing or breaking her phone (RPL = Random Phone Loss).

havior of individuals with friends who are more likely to lose or break their phone, or with friends that are more likely to post about it publicly, may be fundamentally different. To address such concerns, we directly control for the number of friends who have posted publicly about losing or breaking their phones in the previous year as well as for the number of friends who have public statuses by default.⁸

While posting about breaking or losing one’s phone leads to a substantial increase in the average probability of obtaining a new phone, there is substantial heterogeneity in the size of this increase across individuals with different characteristics. For example, Panel B of Figure 2 shows that, among individuals who publicly post about losing their phones in week 0, the probability of getting a new phone in that week is 11% for individuals over the age of 30, while it is only about 9% for individuals under 30 years of age. How many friends purchase a phone in a given week is therefore not only affected by *how many* friends lose their phones in that week, but also by *which* friends lose their phones. Under our assumption that phone loss is a conditionally random event, which friends lose their phones is also plausibly random. We use this insight to further improve the power of our instrument.

Specifically, we exploit small-sample variation in whether those friends who randomly lose their phones in a given week are more or less likely to purchase a new phone, conditional on the distribution of this propensity among all friends. For example, one could use the average age among people posting about a random phone loss as an instrument, controlling for the average age among all friends. Many other demographic characteristics are also correlated with a user’s conditional probability of buying a new phone, and all of these characteristics (and their interactions) could serve as instruments. However, using many of these potentially weak instruments would risk overfitting the first stage, therefore biasing our instrumental variables estimates towards the OLS estimates. Since fitting the first stage is a prediction exercise, recent literature suggests using machine learning tools to optimally fit the first stage when there are a large number of possible instruments (e.g., Belloni, Chernozhukov, and Hansen, 2014; Mullainathan and Spiess, 2017; Peysakhovich and Eckles, 2017; Athey, 2018; Chernozhukov et al., 2018). We build on the ideas in this work and use a neural network to create a single propensity score from a large space of possible instruments:⁹

⁸Additionally, it is important that having friends lose or break their phones in a given week is not correlated with individuals losing or breaking their own phones in that week. One reason for such a correlation could be common experiences that are correlated with breaking or losing a phone (e.g., a bachelor party, a trip to the beach, or time spent in a high-crime area). To assess whether phone loss events are temporally correlated across friends, we perform a series of tests on users who post about losing or breaking their phones in week t , calculating the probability that one of their friends posts about losing or breaking their phones in each week from $t - 5$ to $t + 5$. We were unable to find evidence that users lose or break their phones at the same time as their friends (see Appendix A.1.2). Even though such concerns seem to be minor, we include a control indicating whether the user has posted about a random phone loss in all regressions that make use of this instrument.

⁹The intuition behind this instrumental variables approach is similar to that employed in a number of papers that exploit the random assignment of judges or loan officers (e.g., Liberman, Paravisini, and Pathania, 2017; Dobbie, Goldin, and Yang, 2018). The idea in those papers is that some loan officers are more likely to approve a given loan. To the extent that the assignment of loan officers is random, the identity of the loan officer can be used to instrument for the approval decision. In many of these setting, researchers observe multiple decisions by the same loan officer (or judge), allowing them to estimate a loan-officer-specific leniency measure (using leave-one-out approaches to avoid overfitting the first stage). In our setting, we do not observe each individual with sufficient frequency to estimate an individual-specific probability of replacing a randomly-lost phone. However, we can estimate which characteristics of individuals are associated with this probability. Machine learning techniques allow us to determine the optimal functional form of how these characteristics map into the conditional probability of interest. See Belloni, Chernozhukov, and Hansen (2014) for a related application.

$$ProbBuy\widehat{RandomPhoneLoss}_{i,t} = Prob(\mathbb{1}(BuysPhone_{i,t})|X_{i,t}, \mathbb{1}(RandomPhoneLoss_{i,t}) = 1). \quad (1)$$

The vector $X_{i,t}$ collects a large number of observable characteristics of user i at time t .¹⁰ We train the neural network using data from a separate sample of weeks, 2016-15 to 2016-17 and 2016-23 to 2016-25. This approach, which is similar to the jack-knife IV approach in Angrist, Imbens, and Krueger (1999), allows us to avoid overfitting in-sample noise, thus ensuring that we obtain unbiased estimates when building our instruments based on $ProbBuy\widehat{RandomPhoneLoss}_{i,t}$. Appendix A.2 provides details on the design and the performance of the neural network used to estimate the propensity score.

We then construct the first instrument for the number of friends of person i who purchase a phone in week t by summing these propensities among user i 's friends who post about a random phone loss:

$$Instrument_{i,t}^{Lose} = \sum_{j \in Fr(i)} \mathbb{1}(RandomPhoneLoss_{j,t}) \cdot ProbBuy\widehat{RandomPhoneLoss}_{j,t}, \quad (2)$$

where $Fr(i)$ is the set of all users who are friends with user i . As discussed above, we add controls for the average of $ProbBuy\widehat{RandomPhoneLoss}_{j,t}$ among all of a user's friends in the IV regressions with this instrument. This step allows us to exploit small-sample variation in the probability of replacing a lost phone of the friends who randomly lose their phones in a given week, without capturing a possible direct relationship between the average conditional probability among a user's friends and that user's own probability of purchasing a new phone in that week.¹¹

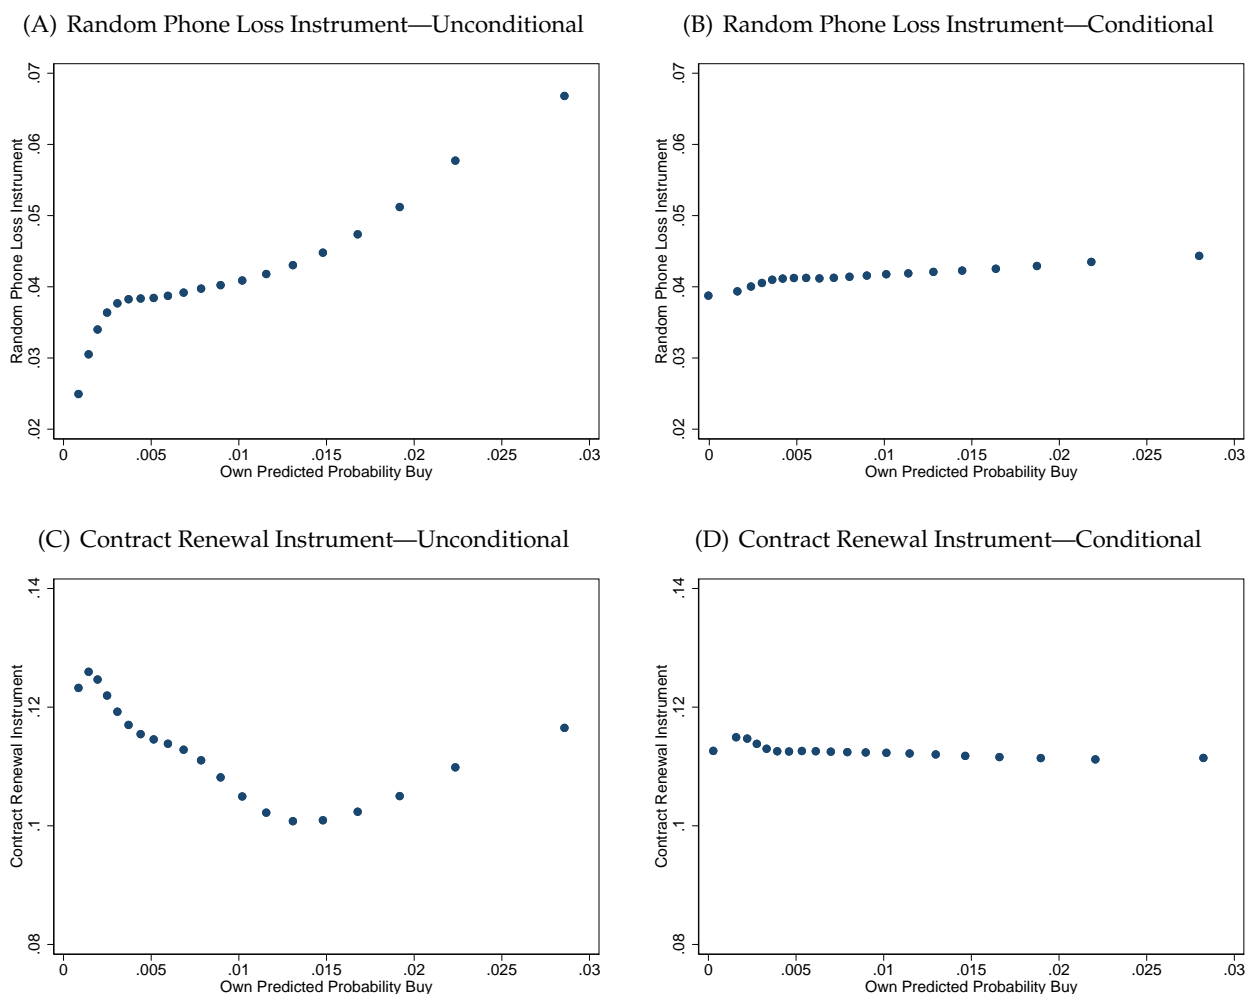
While the exclusion restriction is inherently untestable, we verify its plausibility by exploring whether our instrument is conditionally related to important observable user characteristics. In particular, for each user, we first calculate the unconditional probability that she purchases a phone in a given week, $Prob\widehat{Uncond}_{i,t}$, based on observable characteristics of the user (see Appendix A.2 for details). In Figure 3, we then show the correlation between our instruments and the predicted probability that the user purchases a phone in a given week. In the first row, we explore the random phone loss instrument (equation 2). Panel A shows unconditional relationships. We find that $Instrument_{i,t}^{Lose}$ correlates with a user's own predicted probability of buying a new phone—probably due to homophily. However, Panel B shows that after controlling for the characteristics of a user's overall group of friends

¹⁰We use the following characteristics as features when training our models: current phone age, current phone model, carrier, user age, user browser, Instagram usage flag, U.S. Census region, number of photos posted in the last 4 weeks, education level, friend count, relationship status, activity flags, account age, profile picture flag, number of friendships initiated, subscription count, gender, and area average income.

¹¹We also explore the possibility that the group of friends who would ever publicly post about a random phone loss is a selected subset of all of a user's total friends. In this case, controlling for the average conditional probability among all of a user's friends may not suffice to eliminate a possible direct relationship between the instrument and the errors in the second stage. To address this possibility, we also control for the average conditional probability of purchasing among a user's friends for whom $\mathbb{1}(RandomPhoneLoss_{i,t}) = 1$ at any point in the year prior to our sample period. In the case of a user having no such friends, we set their average probability to a value outside the normal range of the data (in our case, to -1), and we include a binary control for missing data. This procedure allows us to avoid dropping observations when the user had no friends who had $\mathbb{1}(RandomPhoneLoss_{i,t}) = 1$ in the prior twelve months.

(which are also included as controls in our IV specifications), there is no residual relationship between $Instrument_{i,t}^{Lose}$ and the estimated probability that an individual herself purchases a new phone. This lack of conditional correlation between our instrument and observable user characteristics that influence purchasing decisions supports the credibility of our identifying assumption that no such correlation exists on unobservables, either.

Figure 3: Conditional Independence of Baseline Instruments



Note: Panel A shows the unconditional relationship between a user’s own predicted probability to buy a new phone, $ProbUncond_{i,t}$, on the horizontal axis and the random phone loss instrument, $Instrument_{i,t}^{Lose}$, on the vertical axis. Panel B shows the same relationship but conditions on the controls included in Equation 5, with the exception of $ProbBuyRandomPhoneLoss_{i,t}$, the horizontal axis variable. Panels C and D in the bottom row show the analogous relationships for the contract renewal instrument.

2.2 Phone Age Instrument

Our second instrument is based on the fact that during the period of our study, there were two main contract structures in the U.S. cell phone market. The first involved month-to-month contracts in which

a user would purchase her own phone. This type of contract was offered primarily by T-Mobile, AT&T, and MetroPCS. The second contract structure involved carriers subsidizing customers' phone purchases in exchange for a two-year service commitment at a set price. Service of this kind was offered primarily by Sprint and Verizon during that time.

Figure 4 shows the weekly probability of a user obtaining a new phone by the age of their current phone. Panel A shows that this probability generally increasing in phone age, but it spikes when phones cross the two-year age threshold (the dark grey area). Panel B highlights that this spike is concentrated among customers whose service is provided by Verizon or Sprint. As before, we use a neural network to estimate, for each consumer, the probability of buying a new phone in the week when his current phone is two years old, $\widehat{ProbBuy2y}_{i,t} = Prob(\mathbb{1}(BuysPhone_{i,t})|X_{i,t}, \mathbb{1}(Phone2yOld_{i,t}) = 1)$, where $\mathbb{1}(Phone2yOld_{i,t}) = 1$ is an indicator that is set to one for individuals whose phones are between 721 and 735 days old. As suggested by Panel B of Figure 4, a key predictor here is a user's current carrier, but other demographic characteristics also influence this conditional probability. We then instrument for the number of friends who get a new phone with the sum of $\widehat{ProbBuy2y}_{j,t}$ across all friends who are at the two-year phone age threshold in a given week:

$$Instrument_{i,t}^{2y} = \sum_{j \in Fr(i)} \mathbb{1}(Phone2yOld_{j,t}) \cdot \widehat{ProbBuy2y}_{j,t}. \quad (3)$$

Since individuals who have more friends with older phones are plausibly different from individuals with friends who have younger phones on average, we directly control for the number of friends whose phones are between 721 and 735 days old. We also add controls for the number of friends who were at the two-year phone age threshold in the twelve months prior to our sample, as well as the average value of $\widehat{ProbBuy2y}_{j,t}$ among those people, in addition to the average value of $\widehat{ProbBuy2y}_{j,t}$ among all friends. By including these controls, we are effectively using only small-sample variation in the conditional probabilities of a user's friends who are at the contract renewal threshold in a given week, without using variation in the number of these friends. The bottom row of Figure 3 shows that after including these controls, there is no relationship between $Instrument_{i,t}^{2y}$ and a user's own estimated probability of purchasing a phone in a given week, $\widehat{ProbUncond}_{i,t}$.

2.3 Empirical Specification

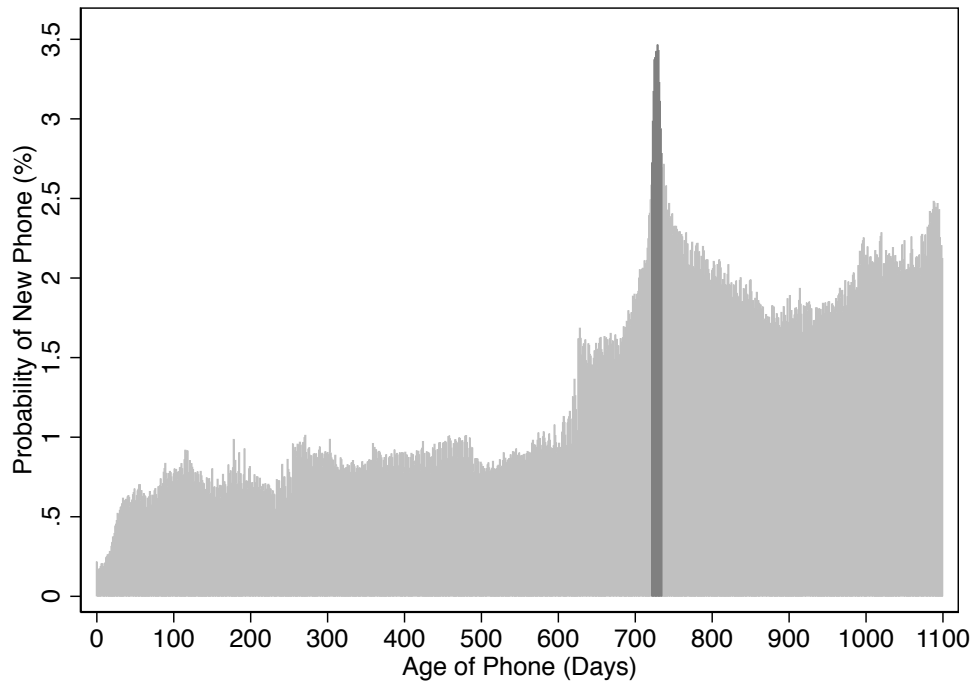
Using these instruments, we implement an instrumental variables (IV) research design to estimate the magnitude of peer effects in the market for cell phones. The first and second stages of the IV regression, respectively, are as follows:

$$FriendsBuyPhone_{i,(t-1,t)} = \delta Instrument_{i,t-1} + \omega X_{i,t} + e_{i,t} \quad (4)$$

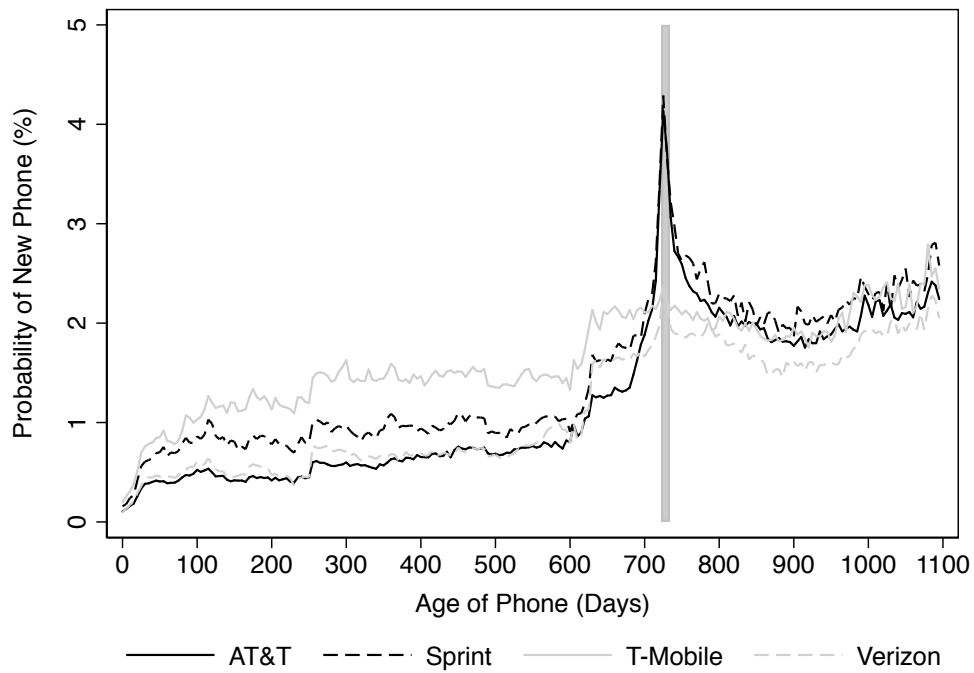
$$\mathbb{1}(BuysPhone_{i,t}) = \beta \widehat{FriendsBuyPhone}_{i,(t-1,t)} + \gamma X_{i,t} + \epsilon_{i,t}. \quad (5)$$

Figure 4: Probability of New Phone by Phone Age

(A) Pooled Probability



(B) Probability Split by Carrier



Note: Panel A shows how a user's probability of getting a new phone varies with the age of their current phone. Panel B shows the same split by user carrier.

The key dependent variable in the second stage, $\mathbb{1}(BuysPhone_{i,t})$, is an indicator of whether individual i purchases a new phone in week t . The vector $X_{i,t}$ represents a rich set of fixed effects and linear controls based on characteristics of the users and their friends. In addition to the controls we already discussed above, we include fully-interacted fixed effects for user characteristics (age bucket \times gender \times education \times state \times week), device characteristics (device \times carrier \times phone age in buckets of 50 days \times week), and friend characteristics (number of friends \times number of friends switching phones in the last 6 months \times week). We also control for the predicted (unconditional) probability that a user purchases a phone in that week, $\widehat{ProbUncond}_{i,t}$.

Our instrument in the first-stage regression is based on shocks to friends in week $t - 1$ (e.g., the number and characteristics of friends who broke their phones in that week). The IV estimate β corresponds to the total user purchases in week t that were induced by the instrument, scaled by the first-stage estimate δ of how many relevant friend purchases were induced by the instrument. This scaling should account for all friend purchases caused by the instrument that occurred prior to the user’s purchasing decision in week t and that could thus have influenced said purchasing decision. As mentioned above, our data do not allow us to precisely pinpoint the timing of purchases, and Figure 2 shows that friends who randomly lose their phones in week $t - 1$ have a somewhat elevated purchasing probability in week t . An analogous, though weaker, increase in purchasing in week t occurs when a user reaches the contract renewal threshold in week $t - 1$. We therefore include all friend purchases in weeks t and $t - 1$ in our endogenous variable, $FriendsBuyPhone_{i,(t-1,t)}$:

$$FriendsBuyPhone_{i,(t-1,t)} = \sum_{j \in Fr(i)} \mathbb{1}(BuysPhone)_{j,t-1} + \sum_{j \in Fr(i)} \mathbb{1}(BuysPhone)_{j,t}. \quad (6)$$

This approach potentially overcounts the relevant number of instrument-induced purchases of new phones by friends, since it can include some friend purchases in week t that occurred after the user has already purchased a phone in that week; as a result, the second-stage coefficient estimates of β provide a conservative measure of the magnitude of peer effects.¹²

3 Peer Effects in Phone Purchasing

We next explore how a user’s propensity to purchase a new phone is affected by the phone purchases of her friends. Section 3.1 presents the baseline estimates of peer effects. Section 3.2 explores the timing of these peer effects, showing that an individual acquiring a new phone increases the aggregate propensity that his friends purchase a new phone for an extended period. In Section 3.3, we explore heterogeneities in both influence and susceptibility to influence across individuals; for example, we highlight that individuals who are more price-sensitive have larger peer effects on their friends. We also discuss the implications of this finding for optimal price-setting strategies of firms.

¹²Using only friend purchases in week $t - 1$ as the endogenous variable would instead undercount the relevant friend purchases induced by the instrument, since it would miss purchases that occurred early in week t (before the user’s own purchasing decision in that week). It would thus understate the first stage (and overstate the second stage), providing an upper bound on the magnitude of peer effects rather than a lower bound as our baseline specification does.

3.1 Baseline Results

Column 1 of Table 2 presents OLS estimates from regression 5. The results suggest that having one more friend purchase a phone in weeks t or $t - 1$ increases a person’s own propensity to purchase a phone in week t by about 0.034 percentage points. This estimate is large relative to a baseline probability of purchasing a new phone of just under one percentage point per week. However, as discussed above, this estimate might also pick up the effects of common shocks or common preference in addition to any peer effects. The rest of Table 2 therefore presents causal effects from instrumental variables estimations. Specifically, columns 2 and 3 show the reduced forms from the random phone loss instrument and the contract renewal instrument, respectively, while columns 4 and 5 show the corresponding second-stage estimates.

Table 2: All Instruments—All Phones

	OLS	Reduced Form		Second Stage	
	(1)	(2)	(3)	(4)	(5)
Friends Buy Phone	0.034*** (0.000)			0.041*** (0.005)	0.026** (0.013)
Instrument		0.048*** (0.006)	0.027** (0.014)		
Controls + Fixed Effects	Y	Y	Y	Y	Y
Instrument Used		Random Phone Loss	Contract Renewal	Random Phone Loss	Contract Renewal
Mean Dependent Variable	0.95	0.95	0.95	0.95	0.95
Number of Observations	335m	335m	335m	335m	335m
F-Statistic Instrument				339,156	55,592

Note: Table shows estimates of regression 5. Column 1 presents the OLS estimate, columns 2 and 3 present reduced form estimates using our two instruments, and columns 4 and 5 present the second-stage IV estimates using the same two instruments. The dependent variable in all specifications is an indicator for whether user i purchases a new phone in week t . We include interacted fixed effects for individual i ’s demographics (age bucket \times state \times gender \times education), individual i ’s beginning-of-week device (current phone \times current phone age in buckets of 50 days \times carrier) and individual i ’s friends (total friends \times number of friends switching phones in the previous 6 months). We control linearly for the user’s unconditional probability of buying a new phone, estimated as described in Appendix A.2 and for the average conditional purchase probability among a user’s friends. In columns 2 and 4, we additionally control for individual and friend posting behavior (the number of friends with public statuses, the number of friends posting in a given week, the number of friends who post about random phone loss in the twelve months prior to our sample, the average conditional probability of buying a new phone among friends who posted about random phone loss in the prior twelve months, and a dummy for whether the user herself posted about a random phone loss in the given week). In columns 3 and 5, we additionally control linearly for the number of friends whose phones are between 721 and 735 days old, the number of friends who have had phones of this age in the twelve months prior to our sample, and the average conditional probability of buying a new phone among those friends. Standard errors are clustered at the individual level. Significance levels: * ($p < 0.10$), ** ($p < 0.05$), *** ($p < 0.01$).

Both second-stage IV estimates are similar in magnitude to the OLS estimate: the IV estimate is slightly larger than the OLS estimate when using the random phone loss instrument, and it is slightly smaller than the OLS estimate when using the contract renewal instrument; neither of these differences is statistically significant.¹³ This similarity in estimated peer effects across OLS and IV specifications is

¹³The upper-bound IV estimates discussed in footnote 12 are somewhat larger. Our estimate using the random phone loss instrument is 0.061 (SE 0.008), while the estimate using the contract renewal instrument is 0.044 (SE 0.022).

perhaps surprising; one might have expected that common shocks or common preferences would lead to a substantial upward bias in the OLS estimates. In contrast, our result here suggests that—after controlling for observable characteristics of individuals and their friends—correlated unobservable shocks or preferences induce at most a small bias to our OLS estimates, at least when estimating the effect of peer purchases on the near-contemporaneous purchasing behavior of individuals.

The difference in magnitudes across the two IV estimates also highlights that the local average treatment effects (LATEs) we capture using each of these instruments may differ from the average effect in the population. For example, our first instrument captures the average peer effects of individuals who post publicly about losing their phones (and who then immediately purchase a new one) on those individuals’ friends. Our finding suggests that the peer effects exerted by these individuals are somewhat larger than the average peer effects in the population, perhaps because individuals who immediately replace a (partially) broken phone care a lot about phones, and are therefore more likely to influence their friends.¹⁴ In contrast, the IV coefficient using the second instrument identifies the average peer effects from individuals who keep the same phone for two years before replacing it. As can be seen from Table 1, a two-year-old phone is in the right tail of the phone age distribution. This result suggests that users who wait that long to replace their phones may be less interested in up-to-date technology than the average user in our sample, perhaps explaining why eventual purchases by these individuals have a below-average effect on the purchasing behavior of their peers.

These differences in local average treatment effects suggest the presence of substantial heterogeneities in peer effects, both along characteristics of the potential influencers and characteristics of the individuals who are potentially influenced.¹⁵ We explore these heterogeneities, which have important implications for firms’ marketing strategies and price-setting behaviors, in Section 3.3.

3.2 Peer Effects at Longer Horizons

The baseline specification in the previous section explores the effects on an individual’s phone-purchasing behavior immediately following a new phone acquisition by a peer. In this section, we explore two related questions. Firstly, for how long does the purchase of a phone by a peer influence an individual’s own purchasing behavior? Secondly, do these peer effects primarily represent a retiming of already-planned purchases, or do they generate purchases that would not have happened otherwise?

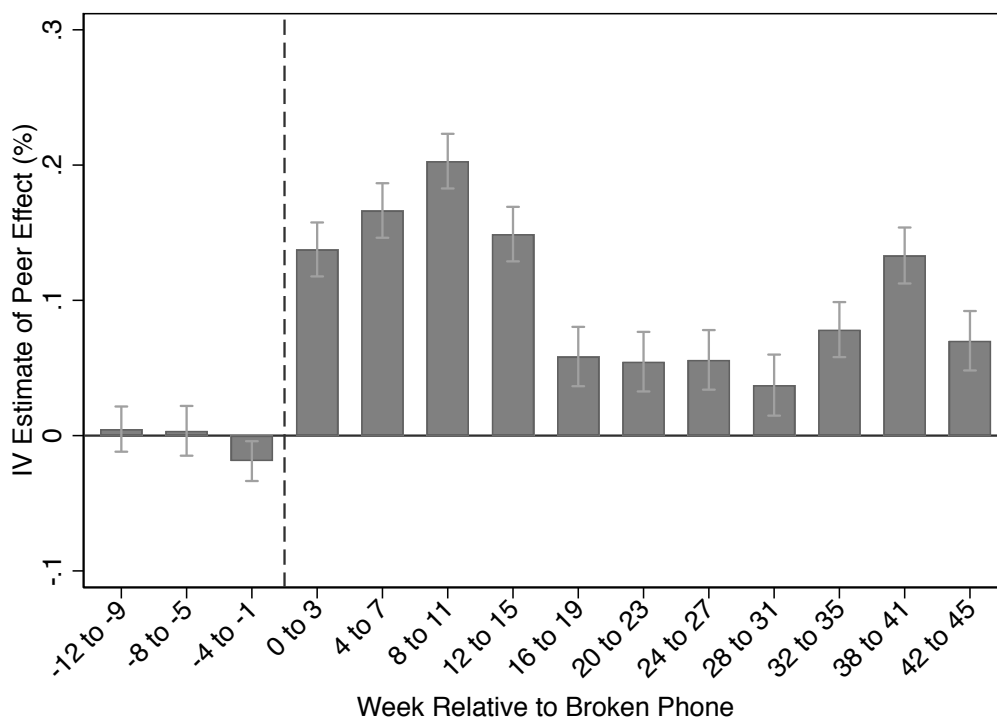
To address these questions, we expand the horizon over which we measure our users’ phone purchasing behavior to include up to 45 weeks following the initial phone purchase by a peer. Specifically, we construct dependent variables of the form $\mathbb{1}(BuysPhone_{i,(t,t+3)})$, $\mathbb{1}(BuysPhone_{i,(t+4,t+7)})$, and so on,

¹⁴In addition, due to homophily, the users who are friends with these people may themselves be more interested in phones, so their own purchasing behavior may be more affected by peer effects than that of the average person.

¹⁵The differences in LATEs across instruments also suggest a potential alternative interpretation of the observation that OLS and IV estimates have similar magnitudes. In particular, it could still be the case that the OLS estimate presents a substantially upward-biased estimate of the true average peer effect in the population, and at the same time that the IV estimates both correspond to LATEs capturing the peer effects from relatively influential individuals, with the two effects approximately offsetting each other. Since it is impossible to distinguish between these interpretations, we will focus on interpreting IV estimates wherever possible.

to capture whether the user purchases a new phone during a number of four-week periods. In Figure 5, we report the β -coefficients found when using these variables as endogenous variables in our baseline regression 5.¹⁶ Though these regressions are similar to our baseline specification, the interpretation of the longer-horizon effects is somewhat more complicated. In particular, since individuals and their friends generally have many friends in common, second-degree peer effects become increasingly relevant at longer time scales: a friend’s purchase in week t may influence a common friend’s purchase in week $t + 1$, which in turn affects the user’s own purchasing decision in week $t + 2$. The coefficients presented in Figure 5 provide the LATE associated with a friend purchasing a new phone in weeks t or $t + 1$ on the user purchasing at various horizons, capturing both the effect of the initial friend purchase and the effect of any purchase by a common friend that was induced by the initial purchase.

Figure 5: Peer Effects at Alternative Horizons



Note: Figure shows estimates of regression 5 at various horizons. The dependent variables are indicator variables for whether a user purchases a new phone in the given four-week period. The IV coefficients capture the effect of friend purchases in weeks $t = 0$ or $t = 1$, induced by a random phone loss in week $t = 0$. Error bars show 95% confidence intervals.

A number of patterns emerge from the instrumental variables coefficients in Figure 5. Firstly, having an extra friend purchase a new phone in response to a random phone loss is not associated with an elevated probability of a user herself purchasing a new phone in the weeks prior to the random phone loss by the friend. This finding provides further support for the exclusion restriction associated with the random phone loss instrument, which requires that individuals with and without a randomly-induced

¹⁶It is not possible to perform a version of this regression using our contract renewal instrument, since this instrument is correlated over time (e.g., a friend whose phone is two years old in week t cannot have a two year-old phone in week $t + 5$).

phone purchase by a friend would behave conditionally similarly in the absence of the random friend purchase. Secondly, Figure 5 shows that the effect on user purchasing of having an extra friend randomly buy a new phone in week $t = 0$ is roughly as large over the first four weeks following the friends' purchase as it is over each of the subsequent 3 months. After that, the aggregate effect declines and stabilizes, with the exception of the weeks around Christmas (weeks 32–35 in Figure 5), when it rises again. During the period that we observe, the aggregate effect on own purchasing behavior in response to a friend replacing a lost phone does not show signs of a reversal. This finding implies that peer effects induce an increase in the total level of phone purchases, and not merely a shift in the timing of a fixed number of purchases.¹⁷

These findings have important implications for the role of peer effects across a number of domains. In particular, our results highlight that the value to a firm of acquiring new customers is substantially larger than the direct contributions of those customers to firm profits, since peer effects can create a long-lasting increase in follow-on demand from the new customers' friends. Our evidence also suggest that there could be similar social multiplier effects from government stimulus programs such as tax rebates aimed at increasing consumer spending. Indeed, our findings imply that empirical strategies that estimate the overall effect of such programs by studying only the consumption response of those individuals who receive a stimulus payment provide a lower bound on overall the aggregate effect of the policies.¹⁸

3.3 Heterogeneities in Treatment Effects

The observation in Section 3.1 that our two instruments identified LATEs of different magnitudes hinted at the presence of substantial heterogeneities in peer effects. In particular, it suggested that those friends whose behavior was shifted by each of our instruments might be differentially influential. To further explore such heterogeneities, we next analyze how peer effects vary with the observable characteristics of users and their friends. These heterogeneities are estimated with IV regressions using the random phone loss instrument, which has the most power; Appendix A.3 provides the exact regression specifications. All patterns of heterogeneity in the resulting LATEs are similar to the patterns of heterogeneity in the corresponding OLS estimates, suggesting that they are a feature not only of the local average treatment effects identified by our random phone loss instrument, but also of the (potentially biased) average treatment effects obtained through OLS analysis.

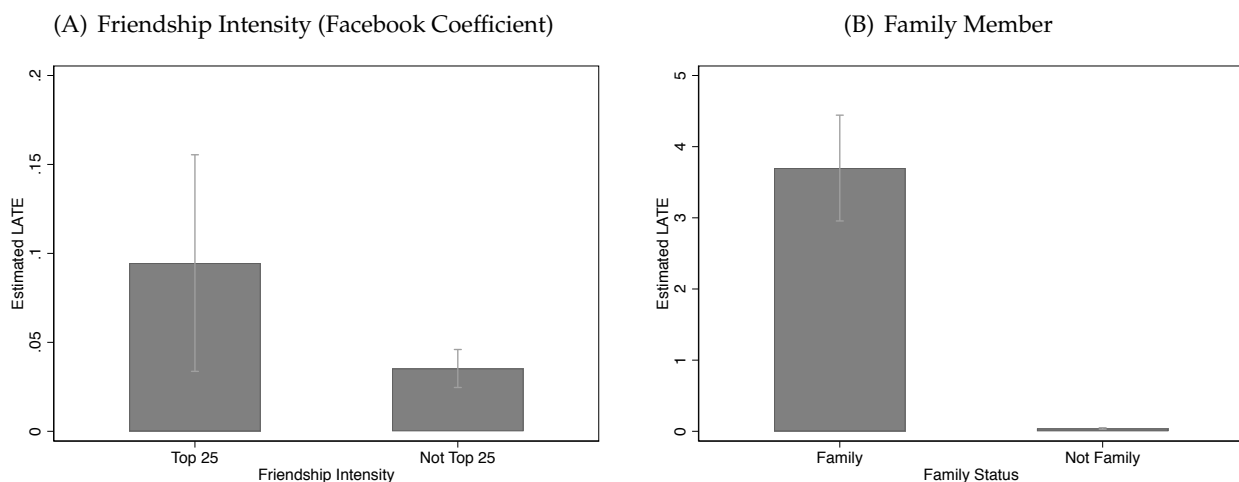
Heterogeneity by Relationship Characteristics. We first explore whether there are some types of

¹⁷This result does not mean that no individuals have their purchases pulled forward through peer effects. Indeed, in all weeks $t' > 1$, there are two countervailing forces that determine the aggregate effect of a random phone purchase in week $t = 0$ and $t = 1$ on the total purchases by all the person's friends. Firstly, there are potentially negative effects on the purchasing probability of people who had their purchases pulled to previous weeks $0 < t < t'$. However, Figure 5 shows that any such effects are more than offset by positive effects on the number of total purchases through delayed or higher-order peer effects.

¹⁸In our application, we are unable to determine whether the additional phone purchases from peer effects come as a result of reduced savings or as a result of switching consumption from other spending categories. Therefore, while we can say that overall phone spending is increasing, we cannot definitely say that overall spending increases—which would likely be the object of interest for evaluating aggregate stimulus policies.

social relationships that generate larger peer effects than others. First, we measure whether peer effects from closer friends are larger than peer effects from less-close friends. To measure the closeness of friendship links, we rank them according to their “Facebook coefficient,” a Facebook-internal model of tie strength based on characteristics such as mutual friends and interaction frequency, similar to Gilbert and Karahalios (2009). Panel A of Figure 6 shows that the estimated peer effect from a friend in the top 25 closest friendships is more than twice as large as the peer effect from a friend who is not in the top 25. It is unsurprising that peer effects from closer friends are larger. Purchases by these friends may be more salient. Further, it is likely that individuals are more willing to trust information that they receive from closer peers. Finally, the desire to keep up with closer friends may be higher than the desire to keep up with friends who are less close.

Figure 6: Peer Effect Heterogeneity by Relationship Characteristics



Note: Figure shows instrumental variables estimates of equation 5 using the random phone loss instrument. Estimated peer effects are split by the characteristics of the relationship between the user and the peer. Panel A splits friends into those inside and outside the top 25 based on a model of friendship intensity; Panel B splits friends based on whether they self-report to be family members. We report the full specifications in the Appendix. Error bars show 95% confidence intervals.

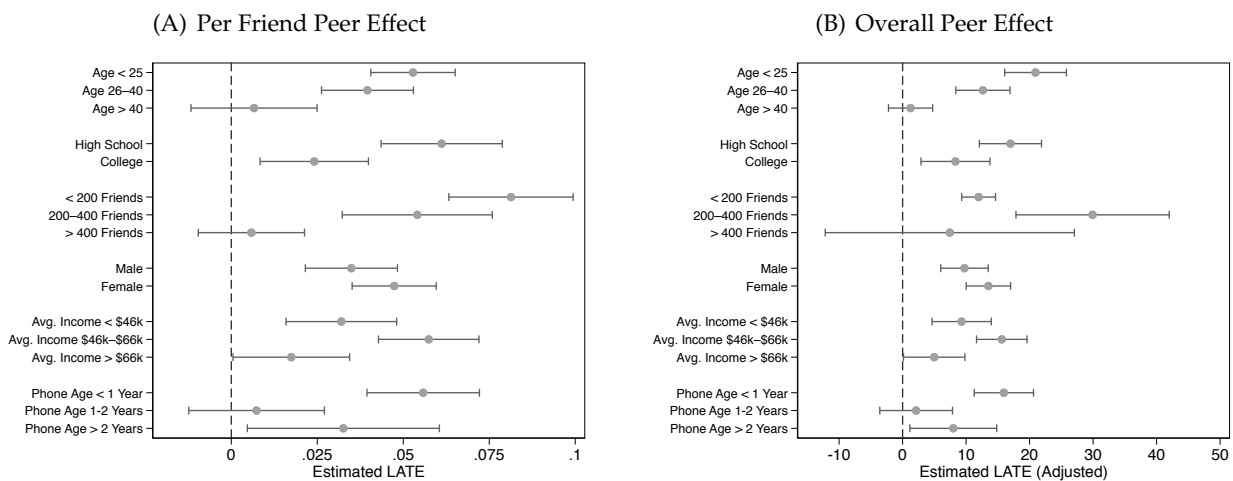
We also analyze whether peer effects from family members are larger than peer effects from other friends. We rely on self-identified family relationships on Facebook to classify friendship links into family and non-family links (only a small proportion of friendship links are identified as a family relation). Panel B of Figure 6 shows that the magnitude of peer effects from family members is about 90 times as large as the peer effects from non-family members, though both are large and statistically significant.¹⁹ There are a number of possible explanations. Firstly, phone purchases by family members are likely to be much more salient than purchases by the average Facebook friend, providing stronger possibilities for social learning. Also, it is possible that family members shop for phones jointly: they

¹⁹Family members are a small enough share of all friends that the estimated “non-family-member” peer effect is similar in magnitude to the average peer effect from all friends. We do not believe that “non-family-member” peer effects correspond primarily to peer effects through family plans from non-identified family members, because such peer effects would primarily lead to correlated purchasing behavior between a peer and a friend that occurs very close in time; in contrast, Section 3.2 showed that much of the peer effects generate user purchases many weeks and months after the peer purchase.

may even share a family plan that will generate incentives for correlated purchasing.²⁰

Heterogeneity by Friend Characteristics. We next explore heterogeneity in the magnitude of peer effects exerted by different individuals. Identifying characteristics of socially influential individuals is an important exercise for marketing researchers and practitioners; “influencer campaigns” are now an integral part of most consumer marketing strategies (see Brown and Hayes, 2008; Ferguson, 2008; Tucker, 2008; Goldenberg et al., 2009; Bakshy et al., 2011; Iyengar, Van den Bulte, and Valente, 2011; Aral and Walker, 2012; Kumar and Mirchandani, 2012). Here, we contribute to this research effort by documenting demographic characteristics that are indicative of large social influence—and by exploring how social influence and price sensitivity are correlated across demographic groups. We discuss that the latter correlation has important implications for firms’ dynamic price-setting behavior.

Figure 7: Peer Effect Heterogeneity by Friend Characteristics



Note: Figure shows instrumental variables estimates of equation 5 using the random phone loss instrument. Estimated peer effects are split by characteristics of the peer with the random phone loss. Panel A shows the mean peer effect a user in each group exerts on each of her friends. Panel B reports the average total influence of a user in each group on all of her friends, computed by multiplying the coefficients found in Panel A with the average number of friends in each demographic group. We report the full specifications in the Appendix. Error bars show 95% confidence intervals.

Figure 7 documents heterogeneity in peer effects along peer demographic characteristics. Panel A shows the “per friend” peer effect, corresponding to the causal effect of a purchase of a new phone by a person with those characteristics on average on each of their friends. Panel B measures the “overall” peer effect, which adjusts the per-friend peer effect by the fact that different demographic groups have differentially many friends. This second category is of particular interest for designing influencer-based marketing campaigns. We find that younger individuals exert larger peer effects on each of their friends. Combined with the fact that these individuals have more friends on average, we find that

²⁰These peer effects do not capture the “passing on” of old devices to family members. There are two reasons. Firstly, the local average treatment effect is identified using the random phone loss instrument, which means that there usually is no working old device that one could pass on. Secondly, we show below that most peer effects are for the same device, suggesting that the influencer and the influenced purchase the same phone.

the overall peer effect exerted by individuals declines substantially in age. This finding suggests that acquiring younger customers is more valuable to firms than acquiring older customers, at least in the phone market, since younger customers will generate more follow-on demand through peer effects.

We also find that the peer effects exerted by individuals who report high school as their highest education level are larger than the peer effects exerted by individuals who report having gone to college. In addition, we find that the per-friend peer effect exerted by individuals is declining in the number of friends that they have, perhaps because the marginal friend is less close. However, despite the declining influence on each friend, the overall peer effects do not follow a similarly monotonic pattern. Users with between 200 and 400 friends seem to have the most influence in aggregate, having a large per-friend peer effect and relatively many friends. The standard error for our estimate of the total influence of users with more than 400 friends is quite large and prevents us from making strong conclusions about that group. We find that women are somewhat more influential than men, although these differences are relatively small. We also find that users from middle-income areas are more influential than users from richer or poorer areas. Users who lose a phone that is less than one year old have a large influence on the purchasing behavior of their friends; individuals who lose a phone that is older than one year have much smaller peer effects on their friends (recall from Table 1 that the median phone age in our sample was 386 days). This result highlights that individuals who do not regularly replace their phones—and who are therefore likely to not value new technology as much—exert at most small peer effects on their friends. The peer effects of these people may be lower both because they are less likely to talk to their friends about having a new phone, and because they may be perceived as less-valuable sources of information when they do talk to their friends.

Peer Influence vs. Price Sensitivity. One implication of peer effects is that the aggregate demand curves faced by firms are more elastic than individual demand curves (see Glaeser, Sacerdote, and Scheinkman, 2003; Becker and Murphy, 2009). The size of this difference depends on the correlation between individuals' price elasticities and the magnitude of the peer effects they exert. Specifically, if a price cut primarily attracts additional demand from individuals who exert only small peer effects, the difference between the individual and aggregate demand curves will be substantially smaller than when a price cut primarily increases the demand of individuals who exert large peer effects.

To explore whether the most influential individuals have relatively high or relatively low price sensitivity, we split individuals into eight mutually-exclusive groups along the interacted dimension of user age (above or below 35 years), user phone age (above or below one year), and user gender. We estimate the per-friend influence and the total influence for each of these eight groups using instrumental variables specifications similar to the ones described above and in Appendix A.3. We also measure the price sensitivity of each group by calculating the percentage increase in the number of users in each group who purchase an iPhone 6s or iPhone 6s Plus in the week before and after a major price cut in September 2016. All estimated price sensitivities are positive, suggesting that all groups increased their average purchasing probability following the price cut. The increase ranges from about

1% to about 15% across groups.²¹ Next, we explore the correlation between peer influence and price sensitivity across these eight groups. We find the correlation to be strongly positive: the correlation with per-friend influence is 0.45, while the correlation with total influence is 0.52.²² This result highlights that price cuts disproportionately attract additional demand from individuals who are relatively influential, suggesting that the deviations between individual and aggregate demand curves in this market are likely to be large. The higher implied price elasticity of aggregate demand will push firms towards setting lower prices than they would in the absence of peer effects. Through this channel, peer effects are a force that lowers markups and improves consumer welfare and allocative efficiency.

The positive correlation between price sensitivity and peer influence can also provide an explanation for the sometimes puzzling observation that many markets clear through queuing rather than through price adjustments. If higher prices disproportionately reduce demand from those individuals with large peer effects on their friends, then an optimal dynamic pricing strategy might be willing to trade off lower revenues today in return for additional sales generated through peer effects in future periods. In other words, while increasing the price would increase revenues today, it might reduce overall long-run revenues due to substantially lower peer effects going forward. In scenarios in which demand exceeds supply, but firms do not want to increase prices to avoid selling only to individuals who exert low peer effects, an alternative assignment mechanism is required. Assignment via queuing is likely to disproportionately select individuals who might exert the largest peer effects among those willing to buy at the low price. This mechanism rationalizes, for example, why Apple does not increase the price for its iPhones, despite the large queues outside its stores around device release dates.²³

In addition to these implications for optimal price setting, the positive correlation between price sensitivity and peer effects also affects the interpretation of various research designs to estimate consumer price elasticities. For example, firms regularly design experiments that offer a random subset of individuals a price discount. The effect of lower prices on consumer demand is then estimated by comparing the purchasing behavior of those individuals who received the discount to the purchasing behavior of those individuals who did not. Such experimental designs will tend to underestimate individual-level price elasticities in the presence of peer effects, since purchases by treated individuals can increase the purchasing probability of the control group through peer effects. This effect will be

²¹On September 7, 2016, Apple announced an immediate price cut of \$100 for the iPhone 6s and the iPhone 6s Plus. We use purchasing data from one week on each side of this date to measure price sensitivity, but our findings are robust to comparisons that use several weeks on either side of the price cut to determine the price sensitivity of each group.

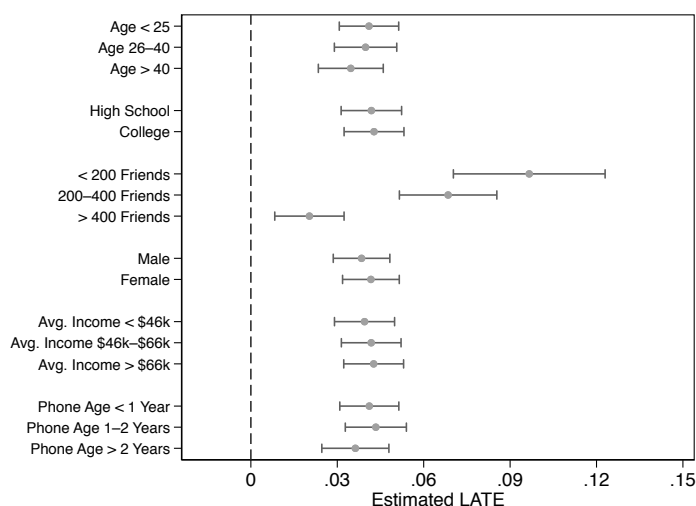
²²We also expand this exercise by further splitting each group into those with more or fewer than 300 friends, providing us with estimates of peer influence and price sensitivity for 16 mutually-exclusive groups. Despite the fact that the estimates for peer influence are substantially noisier, the correlations across these objects are 0.32 and 0.38 for the per-friend and total peer effects, respectively. Running IV regressions with more than 16 endogenous variables is not computationally feasible.

²³A complementary explanation for Apple's decision to not increase prices to clear the iPhone market upon product release is that it generates free marketing when individuals observe the lines, or when media outlets report about these lines. However, this mechanism is unlikely to represent the entire explanation, since similar queuing assignment also occurs for products that are entirely sold online, such as various fashion items, restaurant bookings, or Broadway tickets. For these sales, the characteristic that assigns the goods (which is likely correlated with peer influence) is not the willingness to physically queue outside a store, but rather the willingness to be online at the precise (and often intentionally inconvenient) time that the sales open.

larger when stronger peer effects are exerted by those individuals who respond more to the price cut, as appears to be the case in our setting (see Blake and Coey, 2014, for a related discussion).

Heterogeneity by User Characteristics. Figure 8 explores heterogeneities in the susceptibility to influence of different individuals, separating users along the same demographic characteristics as in Figure 7. There are only small differences in susceptibility to influence across most demographic groups. The exception is that a user’s number of friends is a major determinant of their susceptibility to influence from the average friend—consistent with the marginal friend being less close, and therefore less influential for a user’s purchasing behavior.

Figure 8: Peer Effect Heterogeneity by User Characteristics

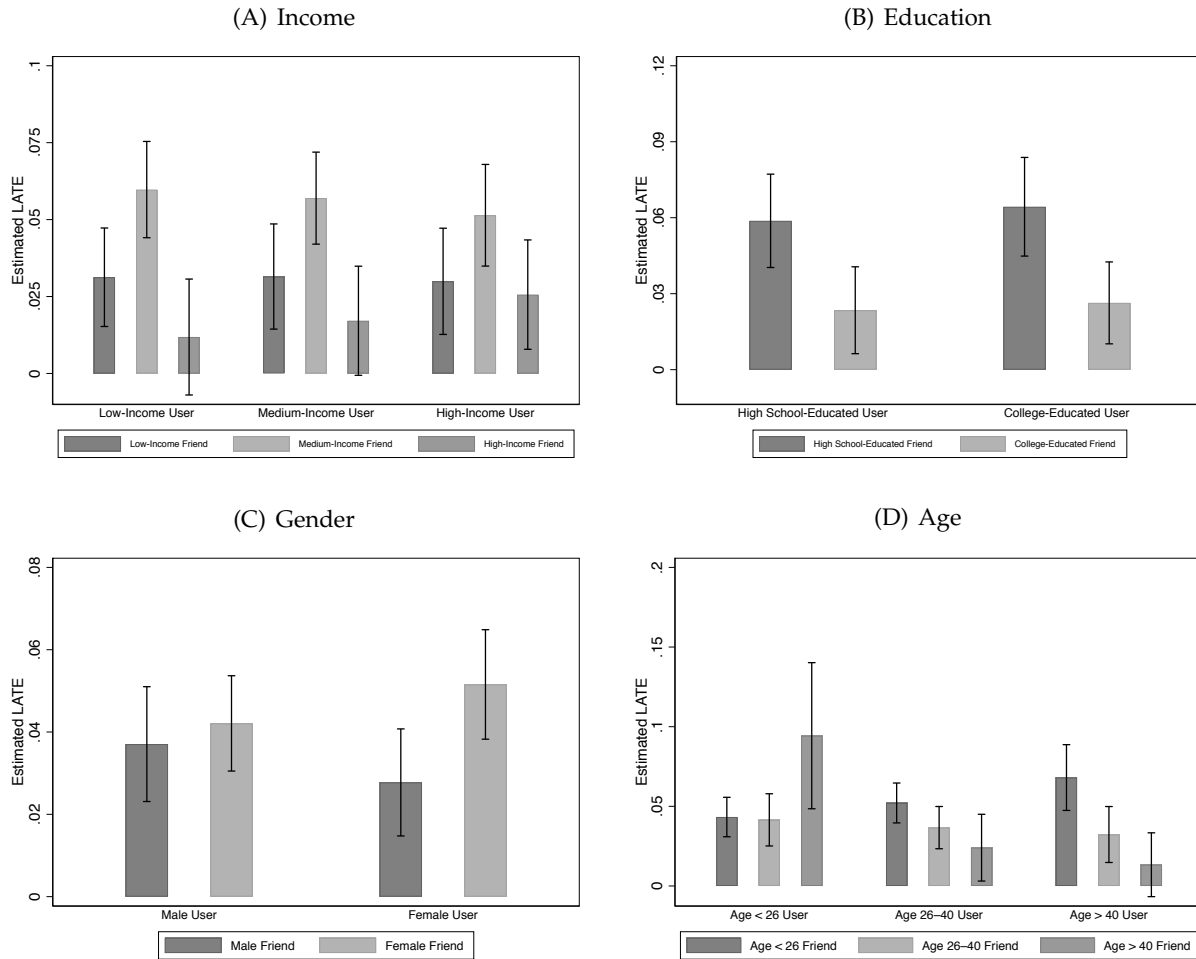


Note: Figure shows instrumental variables estimates of equation 5 using the random phone loss instrument. Estimated peer effects are split by user characteristics. We report the full specifications in the Appendix. Error bars show 95% confidence intervals.

Heterogeneity by Peer and User Characteristics. In the final set of heterogeneity analyses, we explore peer effects along characteristics of both the user and the peer. For example, we explore whether all individuals are primarily influenced by peers who are similar on observable characteristics, or whether all individuals are most influenced by the same types of peers, regardless of their own characteristics.

Panel A of Figure 9 shows the cross-heterogeneity of peer effects by county-level income. Across all user income groups, friends from middle-income areas tend to be the most influential. Panel B shows that, for both high school-educated and college-educated users, high school-educated friends have the largest peer effect. Panel C shows that men and women are both more influenced by female friends than by male friends, though this effect is somewhat larger for female users. Panel D shows that younger users generally have the largest peer effects on their friends, with friends aged 25 years or less having particularly large effects on users older than 40 years. The only exception is the large effect of friends over 40 years old on users below 25 years old, although these peer effect estimates are not very

Figure 9: Peer Effect Heterogeneity by Pairwise Characteristics



Note: Figure shows instrumental variables estimates of equation 5 using the random phone loss instrument. Estimated peer effects are split by user and peer characteristics. We report the full specifications in the Appendix. Error bars show 95% confidence intervals.

precise, and they could be capturing correlated purchasing behavior between parents and children. Overall, these results suggest that individuals who are more influential on average are, in general, more influential on all users, not just those who are similar to them on demographic characteristics.

4 Peer Effects for Specific Phone Purchases

In the previous section, we explored how a user’s decision to purchase any new phone is affected by whether her friends recently acquired a new phone. In this section, we study whether the type of phone that individuals purchase is also influenced by the purchasing behavior of their peers. Specifically, Section 4.1 analyzes whether peer effects in phone purchasing are specific to the brand that the peer purchases, and explores whether there are positive or negative spillovers to competing brands. Section 4.2 further decomposes the observed peer effects to the model level. Our findings highlight the important role of peer effects in shaping the competitive environment in the cell phone market.

4.1 Brand-Level Analysis

We first analyze how the purchase of a specific brand by a peer affects an individual’s own brand-level phone purchasing. We focus on the two major cell phone lines, Apple’s iPhone and Samsung’s Galaxy phones, which are used by 55% and 27% of users in our data, respectively. We pool other brands into a residual category, which includes a variety of phones operating largely on the Android system.²⁴ The set of brand categories we consider is thus given by $C = \{iPhone, Galaxy, Other\}$. We are then interested in understanding how a friend’s purchase of a phone in brand category $c \in C$ affects a user’s probability of buying a phone in the same brand category, as well as their probability of buying phones in a different category. This investigation allows us to answer, for example, whether a friend’s purchase of an iPhone increases a user’s own demand for all phones, including those of iPhone competitor Samsung Galaxy, or whether it primarily pulls demand away from Samsung Galaxys and towards iPhones.

Empirical Approach. As before, OLS specifications are potentially confounded by homophily. To overcome the resultant bias, we adapt the idea behind the random phone loss instrument described above.²⁵ We again argue that the characteristics of the friends who lose their phones in a given week are random conditional on the distribution of these characteristics among all friends. We predict users’ propensities to buy phones in each of our three brand categories. We then use the predicted propensities of the friends who randomly lose their phones as instruments for the number of friends buying a phone of the respective brand, controlling for the predicted propensities among all friends and among friends who lost their phones in the past. Intuitively, we argue that, conditionally upon the characteristics of all of a user’s friends, it is random whether, in a given week, the friends who lose their phones are those who are more likely to purchase iPhones, Galaxys, or other phones.

To implement this research design, we fit a neural network to predict the propensity that individual i will purchase a phone of brand category c in a given week t . We predict both the unconditional probability, $ProbUncond_{i,t}^c$, and the propensity of such a purchase conditional on posting about a random phone loss, $ProbBuyRandomPhoneLoss_{i,t}^c$. We predict these propensities using demographic information about individual i , information about their current cell phone, and information about their friendship network (see Appendix A.2 for details). Many of these characteristics are strong predictors of a user’s likelihood of purchasing a phone of a particular brand. For instance, older users prefer iPhones, while all users are more likely to buy a phone of the same brand as their current device.²⁶ As before, neural networks allow us to uncover non-linear and interactive relationships between these

²⁴Brands in the "Other" category are highly fragmented. In addition to Android phones, the category also includes a small number of BlackBerry and Windows phones, which run on separate operating systems.

²⁵There are two main reasons why we continue to focus on the random phone loss instrument instead of the contract renewal instrument. Firstly, it provides the most powerful shifter of purchasing behavior. Secondly, shocks induced by the random phone loss instrument are uncorrelated over time, allowing us to explore substitution effects over longer horizons.

²⁶We find that 82.9% of current iPhone users purchase an iPhone as their next phone, while 60.8% of current Galaxy users purchase a Galaxy as their next phone. Meanwhile, 54.2% of users with phones from other brands will also purchase a phone other than an iPhone or a Galaxy. Among users of Galaxys and phones in the Other category, those users who purchase a phone in a different category from their current device prefer devices which share their current operating system.

characteristics. For each category c , we then construct our instrument for the number of friends of user i who purchase a phone in category c as:

$$\widehat{FrProbLost}_{i,t}^c = \sum_{j \in Fr(i)} \mathbb{1}(RandomPhoneLoss)_{j,t} \cdot ProbBuyRandomPhoneLoss_{j,t}^c. \quad (7)$$

We argue that this instrument satisfies the exclusion restriction as long as we control for the average conditional probabilities of purchasing phones in each brand category among all friends, which we construct as follows:²⁷

$$FriendAvgProbLost_{i,t}^c = \frac{1}{|Fr(i)|} \sum_{j \in Fr(i)} ProbBuyRandomPhoneLoss_{j,t}^c \quad (8)$$

Figure 10 explores the validity of the exclusion restriction. For each category c , we plot the a user's own $ProbUncond_{i,t}^c$ on the horizontal axis and our instrument for friend purchases of category c , $\widehat{FrProbLost}_{i,t}^c$, on the vertical axis. The left column shows the unconditional relationship and documents that individuals who themselves are more likely to purchase a phone of category c usually have friends that are also more likely to buy such phones, although the relationships are not monotonic. The right column conditions the relationship on a number of control variables also included in our regressions: most importantly, $FriendAvgProbLost_{i,t}^c$. Conditionally on the brand preferences in the overall friend population, the brand preferences of those friends who randomly lose their phones in a given week are essentially uncorrelated with observable characteristics of person i . This finding makes it more plausible that they are also uncorrelated with unobservable characteristics of person i that make that person more likely to buy a phone of category c in that week.

To study peer effects at the brand level, we then perform three instrumental variables regressions: one for each $c'' \in C$. We fit three first stages for each regression, i.e., one for each of the three brand categories $c' \in C$ that a friend could have bought:

$$\begin{aligned} FriendsBuyPhone_{i,(t-1),t}^{c'} &= \sum_{c \in C} \delta_c^{c'} \widehat{FrProbLost}_{i,t-1}^c + \\ &\quad \sum_{c \in C} \phi_c^{c'} FriendAvgProbLost_{i,t}^c + \omega X_{i,t} + e_{i,t}. \end{aligned} \quad (9)$$

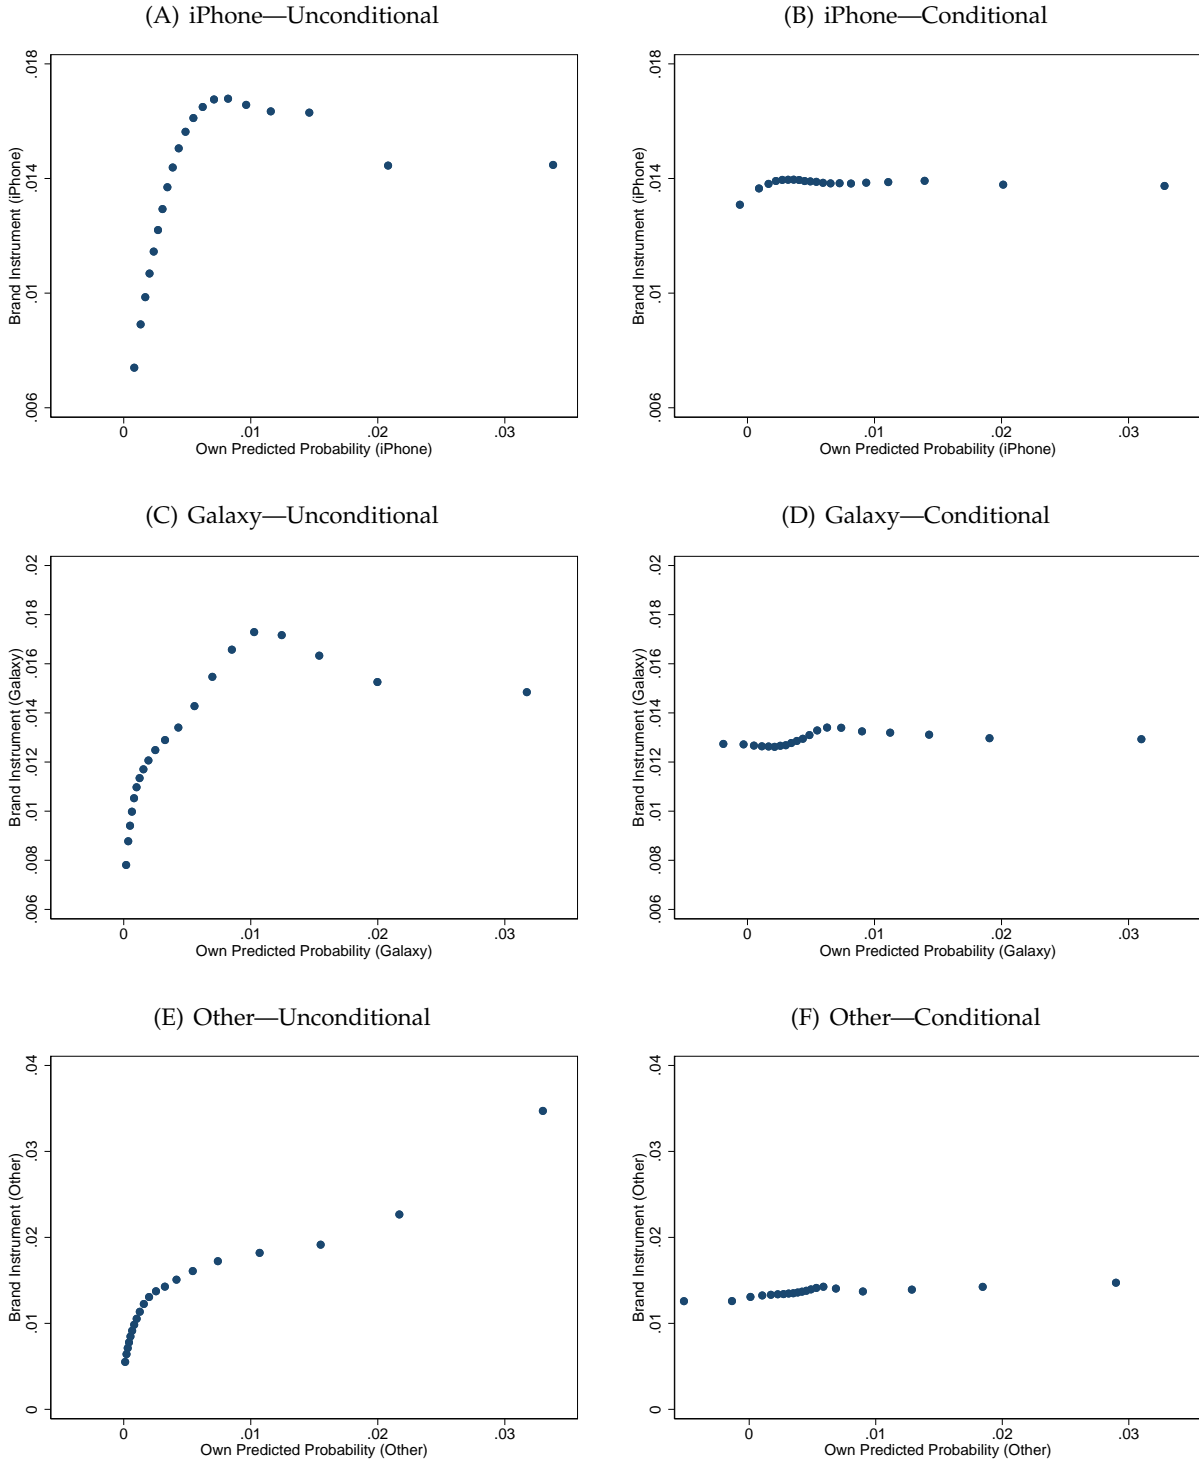
Our three second stages are of the form:

$$\begin{aligned} \mathbb{1}(BuysPhone)_{i,t}^{c''} &= \sum_{c' \in C} \beta_{c'}^{c''} \widehat{FriendsBuyPhone}_{i,(t-1),t}^{c'} + \\ &\quad \sum_{c' \in C} \Phi_{c'}^{c''} FriendAvgProbLost_{i,t}^{c'} + \gamma X_{i,t} + \epsilon_{i,t}. \end{aligned} \quad (10)$$

The indicator variables $\mathbb{1}(BuysPhone)_{i,t}^{c''}$ capture whether person i purchased a phone of brand category

²⁷As described above, it is possible that the sample of users who post about losing or breaking their phones is a selected sub-sample of a user's friends. If this were the case, controlling for the average probability among all friends may not accurately capture the distribution from which the randomly-shocked friends are drawn. We address these concerns by also controlling for the average value of $ProbBuyRandomPhoneLoss_{j,t}^c$ among a user's friends who posted about losing or breaking their phones in the twelve months prior to our sample. Our results are unaffected by the inclusion of these controls.

Figure 10: Conditional Independence of Brand Instruments



Note: Figure shows the relationship between a user's own predicted probability to buy a specific new phone, $ProbUncond_{i,t}^c$, on the horizontal axis and the instrument $FrProbLost_{i,t}^c$ on the vertical axis. The first row shows this relationship for $c = iPhone$, the middle row for $c = Galaxy$, and the bottom row for $c = Other$. The left column shows the unconditional relationship. The right column shows the same relationship but conditions on the controls included in Equation 10, with the exception of $ProbUncond_{i,t}^c$, the horizontal axis variable.

c'' in week t . The coefficients of interest are comprised by the series of $\beta_{c''}^{c''}$, which capture the effects of a friend purchasing a phone in category c' on an individual purchasing a phone in category c'' . The vector $X_{i,t}$ includes the standard controls and fixed effects described in Section 2.3, as well as controls for the unconditional probability that user i buys a phone of each type $c \in C$ in week t and the average of these propensities among the user’s friends. We also estimate a fourth specification with $\mathbb{1}(BuysPhone)_{i,t}$ as the dependent variable, which allows us to examine whether friend purchases of certain brands led to more overall user purchases in week t .

Since some of the (positive or negative) spillovers across brands would likely materialize only over time, we also want to study the effects of a friend purchase on the cumulative probabilities of phone purchases in different brand categories over the subsequent weeks and months. We take an approach similar to that outlined in Section 3.2, constructing dependent variables of the form $\mathbb{1}(BuysPhone)_{i,(t,t+24)}^{c''}$, where c'' is one of our three brands studied. We then perform a second set of instrumental variables regressions of the form outlined in Equation 10, replacing the original dependent variables with these multi-period cumulative purchase indicators.²⁸

Table 3: Peer Effects in Phone Purchasing—Category-Level Analysis

	Dependent Variable: Buys between t and t+1 (%)				Dependent Variable: Buys between t and t+24 (%)			
	(1) iPhone	(2) Galaxy	(3) Other	(4) Any Phone	(5) iPhone	(6) Galaxy	(7) Other	(8) Any Phone
Friends Buy iPhone	0.030*** (0.005)	0.000 (0.004)	-0.007* (0.004)	0.022*** (0.007)	0.331*** (0.024)	-0.003 (0.018)	-0.121*** (0.017)	0.207*** (0.033)
Friends Buy Galaxy	-0.004 (0.009)	0.048*** (0.008)	0.026*** (0.008)	0.070*** (0.015)	-0.196*** (0.043)	0.670*** (0.037)	0.403*** (0.036)	0.877*** (0.063)
Friends Buy Other	-0.013** (0.007)	-0.009 (0.007)	0.073*** (0.008)	0.051*** (0.012)	-0.470*** (0.032)	0.081*** (0.030)	1.438*** (0.033)	1.049*** (0.051)
Controls + Fixed Effects	Y	Y	Y	Y	Y	Y	Y	Y
Mean Dependent Variable	0.39	0.30	0.26	0.95	11.74	6.58	5.91	24.23
Number of Observations	335m	335m	335m	335m	335m	335m	335m	335m

Note: Table shows estimates of regression 10. In columns 1–4, the dependent variables measure purchasing probabilities between weeks t and $t + 1$; in columns 5–8, the dependent variables measure cumulative purchasing probabilities between weeks t and $t + 24$. We include interacted fixed effects for individual i ’s demographics (age bucket \times state \times gender \times education), individual i ’s device (current phone \times current phone age in buckets of 50 days \times carrier) and for individual i ’s friends (total friends \times number of friends switching phones in the previous 6 months). We control linearly for the user’s unconditional probability of buying a new phone in each category c , and for the average conditional and unconditional probability of purchasing a phone in each category among the user’s friends. We additionally control for individual and friend posting behavior (the number of friends with public statuses, the number of friends posting in a given week, the number of friends who post about random phone loss in the twelve months prior to our sample, the average conditional probability of buying a phone of each type c among friends who posted in the prior twelve months, and a dummy for whether the user herself posted about a random phone loss in the given week). Standard errors are clustered at the individual level. Significance levels: * ($p < 0.10$), ** ($p < 0.05$), *** ($p < 0.01$).

²⁸As discussed in Section 3.2, the presence of higher-order peer effects changes the interpretation of the coefficient estimates in the longer-horizon regressions. In these regressions, a larger portion of the measured peer effect will be due to second-order effects, so the coefficient estimated can be considered the “total” peer effect caused by a friend purchasing a phone of brand c at time $t - 1$ or t , including the higher-order peer effects generated by their purchase.

Estimates of Brand-Level Peer Effects. Table 3 shows results from regression 10. Columns 1–4 analyze a user’s purchasing behavior in the week after the friend’s random phone loss, analogous to the baseline specification in Table 2, while columns 5–8 analyze the cumulative purchasing behavior in the 24 weeks following the friends’ random phone loss, though the patterns are similar when we instead study the cumulative purchasing probabilities over the next 12 or 36 weeks. Columns 1 and 5 show the effects on an individual’s probability of purchasing an iPhone. Columns 2 and 6 display the effects on an individual’s probability of purchasing a Galaxy, while columns 3 and 7 show the effects on an individual’s probability of purchasing a phone in the “Other” category. Columns 4 and 8 show the effects on the individual’s probability of purchasing any new phone.

We find that friend purchases in each of our three brand categories lead a user to increase their overall probability of purchasing a phone. In all categories, the same-brand peer effects are positive and larger than any across-brand peer effects. For instance, a friend purchasing a Samsung Galaxy primarily increases an individual’s own probability of also purchasing a Galaxy—both in the period immediately following the friend’s purchase and over longer horizons. In terms of magnitude, the same-category peer effects are largest for devices in the “Other” category and are smallest for iPhones. These findings are consistent with a substantial part of the observed peer effects being the result of information acquisition through social learning. During the sample period, iPhones were the most well-established brand, suggesting that acquiring information through peers was probably less important for that brand. In addition, iPhones were the most expensive brand, thus containing the strongest status signal. As a result, we would expect social learning effects to be relatively less important, and “keeping up with the Joneses” effects to be relatively large, for iPhones compared to other brands. The fact that iPhones have smaller same-brand effects than other brands thus suggests an important social learning channel; Section 4.2 provides additional evidence for this hypothesis.

In addition to these large and positive same-brand peer effects, we also find heterogeneous across-brand demand spillovers. Specifically, when considering demand spillovers over long horizons, we find positive spillovers between purchases of Samsung Galaxy phones and phones in the “Other” category. These two brand categories share the Android operating system. This effect is again consistent with an important role played by social learning: while most of the learning from a friend’s phone purchase is about the precise brand bought by the friend, an individual may also learn about features of the Android operating system, making her more likely to buy any type of Android phone.

On the other hand, demand spillovers tend to be negative across brands that use different operating systems. Friend purchases of phones in the Galaxy or Other categories (which largely use the Android operating system) decrease user purchases of iPhones, which use the competing iOS software. Similarly, friend purchases of iPhones tend to have a negative spillover effect to users’ demands for Galaxy phones and phones in the Other category. It is important to note that these demand spillovers across operating systems could have easily been positive. First of all, it could have been that a user who buys a Galaxy causes her friends to desire more expensive phones—of any type, including iPhones—

through a “keeping up” effect. Secondly, positive across-brand spillovers could have emerged, even across competing operating systems, through the salience channel documented in a marketing literature that shows how advertising can increase sales of (non-advertised) options by reminding people of their existence (e.g., Sahni, 2016; Shapiro, 2018; Sinkinson and Starc, 2018). Thirdly, positive demand spillovers to other brands using different operating systems could have resulted from perception transfers across competing brands (see Roehm and Tybout, 2006; Borah and Tellis, 2016, for related work in the marketing literature). Our finding of substantial negative demand spillovers to competing brands using different operating systems therefore helps researchers to understand the implications of peer effects on the competitive dynamics between firms, and to distinguish them from the spillover effects of marketing activities.²⁹

Summary of Brand-Level Findings. There are four key take-aways from the cross-brand analysis. First, for all three brand categories, there exist large and positive peer effects for same-brand purchases. Second, these same-brand peer effects are largest for the lesser-known but cheaper phones in the “Other” category, and they are smallest for the expensive and well-known iPhones. Third, we find positive different-brand demand spillovers for brands sharing an operating system, and we find negative different-brand spillovers for brands on competing operating systems. Fourth, positive different-brand, same-operating-system spillovers are smaller than the positive same-brand effects. These findings point towards social learning as a substantial contributor to the observed peer effects: when a friend purchases a new phone, individuals learn something about that phone, and about other phones using the same operating system. As a result, demand should increase the most for the specific phone purchased by the friend, and it should increase somewhat less for competing phones that share the same operating system. The importance of this social learning is largest for the least-well-known brands. Some of the incremental same-brand purchases from peer effects correspond to newly-generated demand, and some correspond to a shifting of demand from other brands on competing operating systems.

4.2 Phone Model-Level Analysis

In this final section, we explore peer effects at the model level. Specifically, we analyze whether having a friend buy an iPhone 6s primarily increases a person’s own probability of also purchasing an iPhone 6s, or whether it increases the individual’s probability of purchasing an iPhone in general. To do so, we explore purchases of each of the 20 most common models that individuals switch to in our sample, in addition to three residual categories. The full set of phones P that we study includes iPhone 6s, iPhone 6s Plus, iPhone SE, iPhone 6, iPhone 5s, iPhone 6 Plus, Other iPhone, Galaxy S7, Galaxy S7 Edge, Galaxy Core Prime, Galaxy Note 5, Galaxy J7, Galaxy Grand Prime, Galaxy S5, Galaxy S6, Other

²⁹It is possible that across-brand spillovers were positive even across operating systems during other periods. For example, it is conceivable that in the early years of the iPhone, having friends buy an iPhone helped to establish smartphones as a category, therefore also increasing the sales of non-iPhone smartphones (see Agarwal and Bayus, 2002, for an analysis of new category establishment).

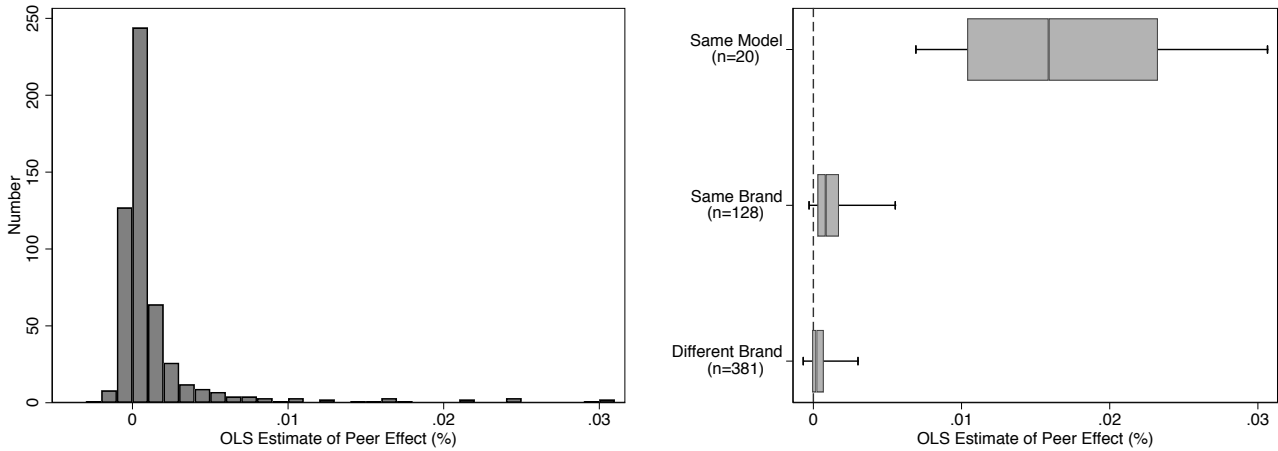
Galaxy, LG Tribute 5, LG K10, LG G5, LG G Stylo, HTC Desire 626s, Alcatel One Touch, and Other.

For this analysis, we cannot use an instrumental variables research design. In particular, while observable characteristics allow us to predict relatively well whether a given individual would purchase an iPhone or a Galaxy, it is much harder to predict whether an individual would buy an iPhone 6 or an iPhone 6s, and the resulting estimated probabilities are highly collinear.³⁰ Instead of an instrumental variables specification, we therefore run the following OLS regressions for each model $p' \in P$:³¹

$$\mathbb{1}(BuysPhone)_{i,t}^{p'} = \sum_{p \in P} \beta_p^{p'} FriendsBuyPhone_{i,t-1}^p + \sum_{p \in P} \pi_p^{p'} \widehat{ProbUncond}_{i,t}^p + \gamma X_{i,t} + \epsilon_{i,t} \quad (11)$$

The coefficients of interest are the parameters $\beta_p^{p'}$, of which there are a total of $|P| \times |P| = 529$. The left panel of Figure 11 shows a histogram of the estimated coefficients. Around 75% of the coefficients are positive. The similarity of the OLS and IV estimates in Section 3.1 suggests that biased coefficients due to correlated shocks may not be a first-order concern for the analysis. In addition, even if our OLS estimates were biased on average, patterns across these coefficients can still be informative about the nature of the underlying peer effects—especially under the assumption that many common shocks should affect these estimates in similar ways.

Figure 11: Specific Model Peer Effects (OLS)



Note: Figure shows the distributions of $\beta_p^{p'}$ coefficients from regression 11. The left panel shows a histogram of the 529 coefficients. The right panel splits these up into three groups. The top bar shows the coefficients when p and p' correspond to the same model; the middle bar corresponds to the coefficients when p and p' correspond to different models of the same brand, and the bottom bar corresponds to the coefficients when p and p' correspond to different brands. The box plots show the 5th, 25th, 50th, 75th and 95th percentiles of the distribution.

The right panel of Figure 11 shows the distributions of $\beta_p^{p'}$ coefficients separately for three groups. The

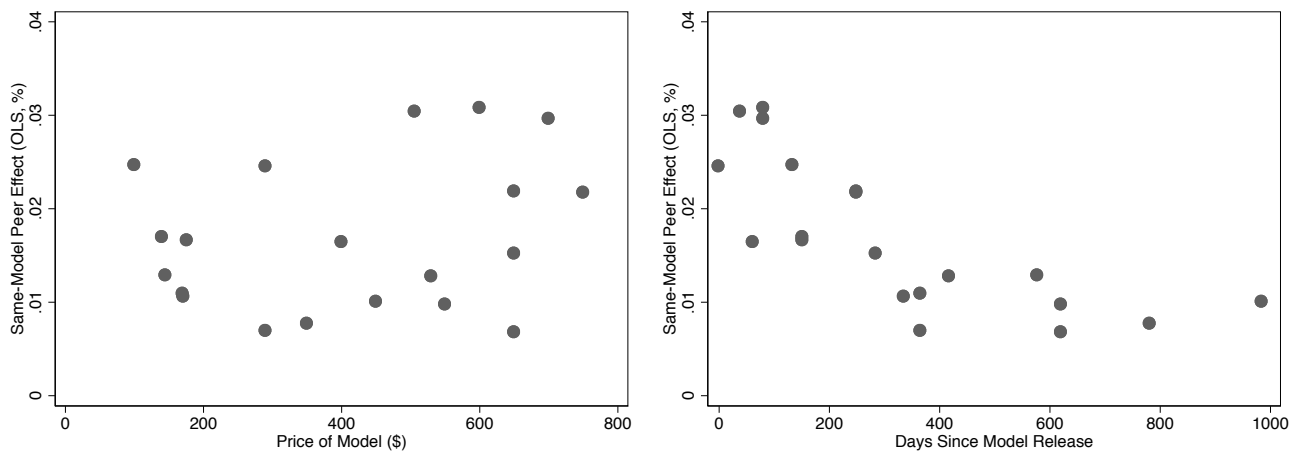
³⁰In Appendix A.2, we describe how we estimate the unconditional probabilities that individuals purchase a phone of each type p in a given week, $ProbUncond_{i,t}^p$. Many of these probabilities are highly correlated within individuals. For example, the predicted propensities to purchase an iPhone 6s and an iPhone 6 have a correlation of 0.89, and the predicted propensities to purchase a Galaxy J7 and a Galaxy S5 have a correlation of 0.90.

³¹ $X_{i,t}$ is a vector of controls analogous to those outlined in Equation 9. Controls for brand-level probabilities are replaced by the corresponding model-level probabilities.

first group includes the 20 same-model coefficients that capture the effects of an individual buying a specific model on the probability of her friends buying the same model. The second group corresponds to same-brand, different-model peer effects. They capture, for example, the effect of an individual buying a Galaxy S7 on his friends' probability of purchasing a different Galaxy model, such as a Galaxy Note 5. The final group includes all peer effects to models from a different brand.³²

By far the largest peer effects are concentrated within the same model. Indeed, the smallest same-model peer effect we estimate is larger than the largest different-model peer effect. There is also substantial heterogeneity across the estimated same-model peer effects. In Figure 12, we plot the 20 same-model $\beta_p^{p'}$ coefficients against the market price of the model during our sample period (left panel), and against the time since the market introduction (right panel). There is no correlation between model price and the estimated same-model peer effect. In other words, low-end models such as the LG Tribute 5 have similarly sized same-model peer effects as more upscale models such as the iPhone 6s Plus. On the other hand, we find that same-model peer effects are substantially larger for more recent models, regardless of the price of these models. As before, these patterns point towards social learning as an important driver for the estimated peer effects (the importance of which should decline as the model becomes more well-known over time). The evidence for a “keeping up” effect, which would likely be more important for more expensive phones, is more limited.

Figure 12: Same-Model Peer Effects



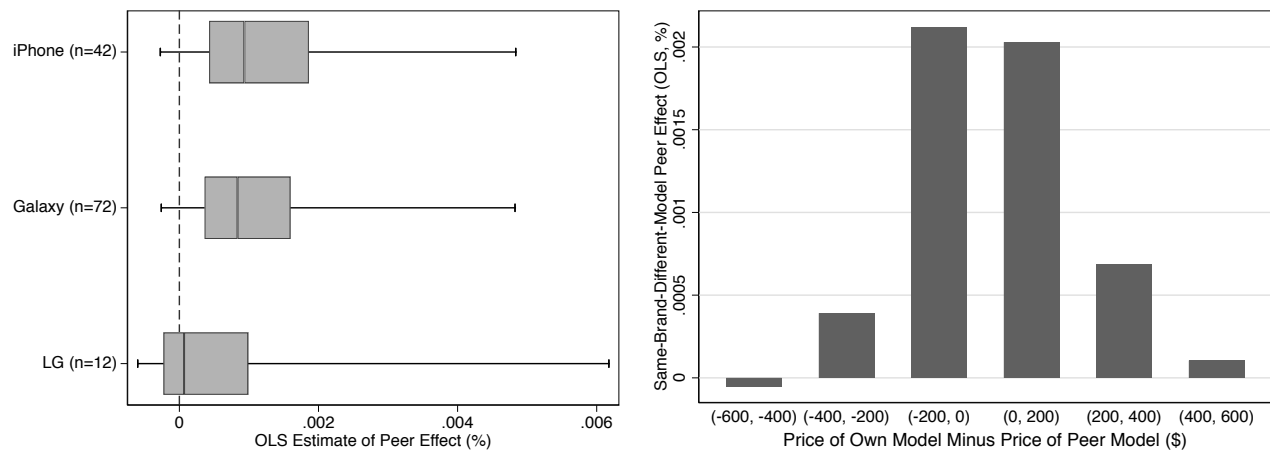
Note: Figure shows scatterplots of the 20 same-model $\beta_p^{p'}$ coefficients from regression 11 against the price of the model (left panel) and the time since the model's release (right panel).

Figure 11 also shows that same-brand, different-model peer effects are more than twice as large as different-brand peer effects, with an average $\beta_p^{p'}$ coefficient of 0.0014 vs. 0.0006. To further explore the

³²We include the coefficients for “Other iPhone” to “Other iPhone” and “Other Galaxy” to “Other Galaxy” in the “Same Brand” category. We include the “Other” to “Other” coefficient in the “Other Brand” category, even though some of these peer effects could still correspond to “Same Model” or “Same Brand” purchases that were so uncommon that they were not split out independently. Among the 20 precise models we split out, there are 3 unique brands with more than one model: iPhone, Galaxy, and LG.

same-brand, different-model peer effects, we split them up according to the three major brands in our data. The left panel of Figure 13 shows that these same-brand, different-model peer effects are largest among the iPhones in our sample and smallest among the LG phones. This finding is consistent with the relatively independent conduct of the marketing campaigns for the LG models (and the model names not indicating any relationship between phones), while both Apple and Samsung tended to jointly market their entire range of phones under a common brand identity. These results suggest that umbrella branding campaigns of different phone models can generate valuable same-brand spillovers through peer effects (see Erdem, 1998; Erdem and Sun, 2002, for analyses of the spillovers of marketing activities due to umbrella branding). The results also highlight additional benefits resulting from line extension strategies beyond the direct effects of advertising spillovers documented, for example, by Balachander and Ghose (2003).

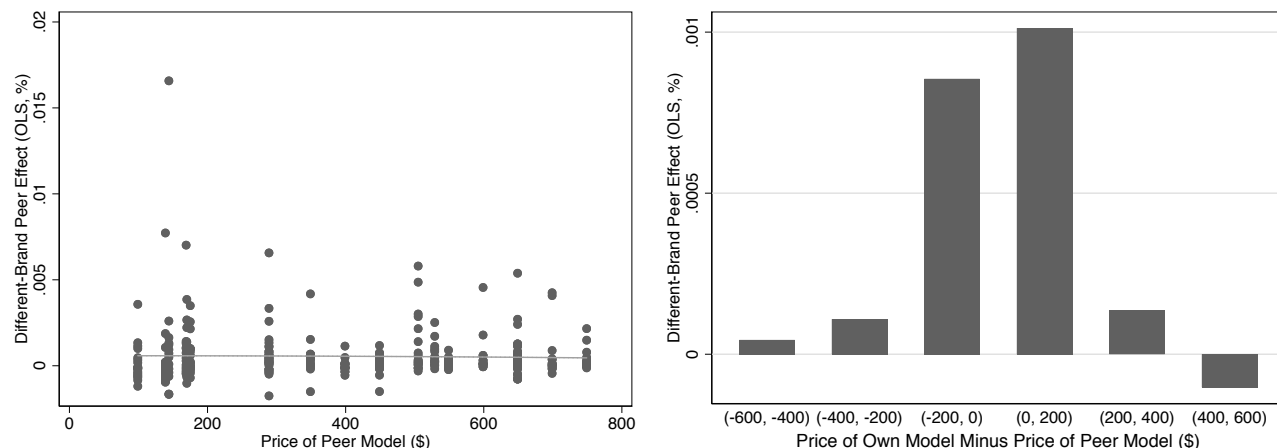
Figure 13: Same-Brand, Different-Model Peer Effects



Note: Figure shows the same-brand, different-model $\beta_p^{p'}$ coefficients from regression 11. The left panel shows the distribution of the effects separately by the three main brands in our sample. The box plots show the 5th, 25th, 50th, 75th, and 95th percentiles of the distribution. In the right panel, we form groups on the basis of the price difference between the models (i.e., $Price(p') - Price(p)$), and plot the average OLS coefficient on the basis of that price difference. Positive numbers capture the peer effects of a friend buying a cheaper model on a person’s probability of buying a more expensive model.

The right panel of Figure 13 shows the same-brand, different-model peer effects split out by the price difference between the two models. Positive numbers on the horizontal axis capture the peer effects of a friend buying a cheaper model on a person’s probability of buying a more expensive model. We find that peer effects are larger for similarly-priced models than they are for models that are either substantially more expensive or substantially cheaper. Importantly, this finding is not just the result of an individual and her friends having similar incomes or being similarly old. Indeed, in all regressions, we directly control for individuals’ estimated unconditional probability of purchasing a certain phone model, in addition to the average of these probabilities across their friends. These findings are consistent with evidence from the marketing literature that cross-product spillovers decrease in magnitude as products become more dissimilar (e.g., Janakiraman, Sismeiro, and Dutta, 2009).

Figure 14: Different-Brand Peer Effects



Note: Figure shows the different-brand $\beta_p^{p'}$ coefficients from regression 11. The left panel shows the distribution of the effects separately by the price of the peer model including a quadratic fit line. In the right panel, we form groups on the basis of the price difference between the models, $Price(p') - Price(p)$, and plot the average OLS coefficient on the basis of their price difference. Positive numbers capture the peer effects of a friend buying a cheaper model on a person's probability of buying a more expensive model.

In the last part of the analysis, we split out different-brand peer effects. The left panel of Figure 14 shows that these peer effects do not vary, on average, with the price of the models. The right panel shows that most of these different-brand peer effects are concentrated on different brand models in a similar price range as the phone purchased by the peer. This is consistent with the patterns and associated interpretations for same-brand peer effects documented above.

Summary of Model-Level Findings. A number of important take-aways result from our model-level analysis. Firstly, same-model peer effects are more than an order of magnitude larger than different-model peer effects. Secondly, these same-model peer effects do not vary with the cost of the model, but rather they are decreasing in the time since the model release, providing additional evidence for an important social learning channel behind the peer effects. Thirdly, same-brand, different-model peer effects are more than twice as large as different-brand peer effects. These effects are largest for brands that co-brand their devices and are also largest within each brand for models of similar value. Finally, across-brand peer effects do not vary with model value, and these effects are largest for models of similar value.

5 Conclusion

In this paper, we document that new phone purchases by friends have substantial, positive, and long-lasting effects on an individual's own demand for phones of the same brand, most of which is concentrated on the particular model purchased by peers. While peer effects expand the overall market for phones, there are substantial negative demand spillovers to competitor brands on different operating systems as a result of a phone purchase by a friend. These negative across-brand demand spillovers

have important implications for firms: losing a customer to a rival firm does not only mean missing out on positive peer effects that this customer could have had, but will also lead to future losses of other customers through competitive peer effects. These findings emphasize how customers' value to a firm exceeds the direct effect that these customers have on the firm's profits. We also find that particularly strong peer effects are exerted by more price-sensitive individuals. This correlation implies a substantial wedge between firms' aggregate demand curves and individual demand curves, and it also highlights the difficulties of learning from individual behavior about aggregate outcomes (and vice versa). While we cannot precisely identify the channel behind the observed peer effects, our results are most consistent with an important role of social learning in explaining the importance of friends' phone purchases on an individual's own purchasing probability.

An interesting question remains about the generalizability of our findings to understanding the decision to adopt products other than cell phones. For example, while the economic insight that the correlation between price sensitivity and peer influence is a central parameter in understanding the role of peer effects for firms, the sign of this correlation does not need to be the same across settings. As a result, we hope that future research will further broaden our understanding of the importance of peer effects in product adoption decisions across a wider range of product categories. In this light, our research emphasizes the increasingly important role of data from online services—such as Facebook, LinkedIn, Twitter, eBay, Mint, Trulia, and Zillow—in overcoming important measurement challenges across the social sciences (see, for example, Baker, 2018; Giglio et al., 2015; Einav et al., 2015; Piazzesi, Schneider, and Stroebel, 2015). Specifically, we hope that the increasing availability of social network data, such as the Social Connectedness Index described in Bailey et al. (2018a), will help to improve our understanding of the effects of social interactions on social, political, financial, and economic outcomes.

References

- Agarwal, Rajshree and Barry L Bayus. 2002. "The market evolution and sales takeoff of product innovations." *Management Science* 48 (8):1024–1041.
- Allcott, Hunt and Judd B Kessler. 2019. "The welfare effects of nudges: A case study of energy use social comparisons." *American Economic Journal: Applied Economics* 11 (1):236–76.
- Angrist, Joshua David, Guido W Imbens, and Alan B Krueger. 1999. "Jackknife instrumental variables estimation." *Journal of Applied Econometrics* 14 (1):57–67.
- Aral, Sinan and Dylan Walker. 2012. "Identifying influential and susceptible members of social networks." *Science* :1215842.
- Athey, Susan. 2018. "The impact of machine learning on economics." In *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.
- Bailey, Michael, Rachel Cao, Theresa Kuchler, Johannes Stroebel, and Arlene Wong. 2018a. "Social Connectedness: Measurement, Determinants, and Effects." *Journal of Economic Perspectives* 32 (3):259–80.
- Bailey, Michael, Ruiqing Cao, Theresa Kuchler, and Johannes Stroebel. 2018b. "The economic effects of social networks: Evidence from the housing market." *Journal of Political Economy* 126 (6):2224–2276.

- Bailey, Michael, Eduardo Dávila, Theresa Kuchler, and Johannes Stroebel. 2017. "House price beliefs and mortgage leverage choice." Tech. rep., National Bureau of Economic Research.
- Baker, Scott R. 2018. "Debt and the response to household income shocks: Validation and application of linked financial account data." *Journal of Political Economy* 126 (4):1504–1557.
- Bakshy, Eytan, Jake M Hofman, Winter A Mason, and Duncan J Watts. 2011. "Everyone's an influencer: quantifying influence on twitter." In *Proceedings of the fourth ACM international conference on Web search and data mining*. ACM, 65–74.
- Bakshy, Eytan, Itamar Rosenn, Cameron Marlow, and Lada Adamic. 2012. "The role of social networks in information diffusion." In *Proceedings of the 21st international conference on World Wide Web*. ACM, 519–528.
- Balachander, Subramanian and Sanjoy Ghose. 2003. "Reciprocal spillover effects: A strategic benefit of brand extensions." *Journal of Marketing* 67 (1):4–13.
- Becker, Gary Stanley and Kevin M Murphy. 2009. *Social economics: Market behavior in a social environment*. Harvard University Press.
- Belloni, Alexandre, Victor Chernozhukov, and Christian Hansen. 2014. "High-dimensional methods and inference on structural and treatment effects." *Journal of Economic Perspectives* 28 (2):29–50.
- Beshears, John, James J Choi, David Laibson, Brigitte C Madrian, and Katherine L Milkman. 2015. "The effect of providing peer information on retirement savings decisions." *The Journal of finance* 70 (3):1161–1201.
- Blake, Thomas and Dominic Coey. 2014. "Why marketplace experimentation is harder than it seems: The role of test-control interference." In *Proceedings of the fifteenth ACM conference on Economics and computation*. ACM, 567–582.
- Bojanowski, Piotr, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2016. "Enriching Word Vectors with Subword Information." *arXiv preprint arXiv:1607.04606* .
- Bollinger, Bryan and Kenneth Gillingham. 2012. "Peer effects in the diffusion of solar photovoltaic panels." *Marketing Science* 31 (6):900–912.
- Borah, Abhishek and Gerard J Tellis. 2016. "Halo (spillover) effects in social media: do product recalls of one brand hurt or help rival brands?" *Journal of Marketing Research* 53 (2):143–160.
- Brown, Duncan and Nick Hayes. 2008. *Influencer marketing: Who really influences your customers?* Routledge.
- Bursztyn, Leonardo, Florian Ederer, Bruno Ferman, and Noam Yuchtman. 2014. "Understanding mechanisms underlying peer effects: Evidence from a field experiment on financial decisions." *Econometrica* 82 (4):1273–1301.
- Chernozhukov, Victor, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. 2018. "Double/debiased machine learning for treatment and structural parameters." *The Econometrics Journal* 21 (1):C1–C68.
- Christakis, Nicholas A and James H Fowler. 2013. "Social contagion theory: examining dynamic social networks and human behavior." *Statistics in medicine* 32 (4):556–577.

- Conley, Timothy G and Christopher R Udry. 2010. "Learning about a new technology: Pineapple in Ghana." *American economic review* 100 (1):35–69.
- De Giorgi, Giacomo, Anders Frederiksen, and Luigi Pistaferri. 2016. "Consumption network effects." Tech. rep., National Bureau of Economic Research.
- DellaVigna, Stefano, John A List, and Ulrike Malmendier. 2012. "Testing for altruism and social pressure in charitable giving." *The quarterly journal of economics* 127 (1):1–56.
- Dobbie, Will, Jacob Goldin, and Crystal S Yang. 2018. "The Effects of Pretrial Detention on Conviction, Future Crime, and Employment: Evidence from Randomly Assigned Judges." *American Economic Review* 108 (2):201–40.
- Duflo, Esther and Emmanuel Saez. 2003. "The role of information and social interactions in retirement plan decisions: Evidence from a randomized experiment." *The Quarterly journal of economics* 118 (3):815–842.
- Einav, Liran, Theresa Kuchler, Jonathan Levin, and Neel Sundaresan. 2015. "Assessing sale strategies in online markets using matched listings." *American Economic Journal: Microeconomics* 7 (2):215–47.
- Erdem, Tülin. 1998. "An empirical analysis of umbrella branding." *Journal of Marketing Research* :339–351.
- Erdem, Tülin and Baohong Sun. 2002. "An empirical investigation of the spillover effects of advertising and sales promotions in umbrella branding." *Journal of Marketing Research* 39 (4):408–420.
- Ferguson, Rick. 2008. "Word of mouth and viral marketing: taking the temperature of the hottest trends in marketing." *Journal of Consumer Marketing* 25 (3):179–182.
- Foster, Andrew D and Mark R Rosenzweig. 1995. "Learning by doing and learning from others: Human capital and technical change in agriculture." *Journal of political Economy* 103 (6):1176–1209.
- . 2010. "Microeconomics of technology adoption." *Annu. Rev. Econ.* 2 (1):395–424.
- Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Andreas Weber. 2015. "Climate change and long-run discount rates: Evidence from real estate." Tech. rep., National Bureau of Economic Research.
- Gilbert, Eric and Karrie Karahalios. 2009. "Predicting tie strength with social media." In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 211–220.
- Glaeser, Edward L, Bruce I Sacerdote, and Jose A Scheinkman. 2003. "The social multiplier." *Journal of the European Economic Association* 1 (2-3):345–353.
- Goldenberg, Jacob, Sangman Han, Donald R Lehmann, and Jae Weon Hong. 2009. "The role of hubs in the adoption process." *Journal of marketing* 73 (2):1–13.
- Han, Bing, David A Hirshleifer, and Johan Walden. 2016. "Visibility Bias in the Transmission of Consumption Beliefs and Undersaving." *Rotman School of Management Working Paper* (2798638).
- Hong, Harrison, Jeffrey D Kubik, and Jeremy C Stein. 2004. "Social interaction and stock-market participation." *The journal of finance* 59 (1):137–163.
- Hoxby, Caroline. 2000. "Peer effects in the classroom: Learning from gender and race variation." Tech. rep., National Bureau of Economic Research.

- Iyengar, Raghuram, Christophe Van den Bulte, and Thomas W Valente. 2011. "Opinion leadership and social contagion in new product diffusion." *Marketing Science* 30 (2):195–212.
- Jackson, Matthew O and Leeat Yariv. 2005. "Diffusion on social networks." In *Economie Publique*. Citeseer.
- Janakiraman, Ramkumar, Catarina Sismeiro, and Shantanu Dutta. 2009. "Perception spillovers across competing brands: A disaggregate model of how and when." *Journal of Marketing Research* 46 (4):467–481.
- Jones, Jason J, Jaime E Settle, Robert M Bond, Christopher J Fariss, Cameron Marlow, and James H Fowler. 2013. "Inferring tie strength from online directed behavior." *PloS ONE* 8 (1):e52168.
- Kim, Yoon. 2014. "Convolutional Neural Networks for Sentence Classification." *arXiv preprint arXiv:1408.5882* .
- Kremer, Michael and Edward Miguel. 2007. "The illusion of sustainability." *The Quarterly journal of economics* 122 (3):1007–1065.
- Kuhn, Peter, Peter Kooreman, Adriaan Soetevent, and Arie Kapteyn. 2011. "The effects of lottery prizes on winners and their neighbors: Evidence from the Dutch postcode lottery." *American Economic Review* 101 (5):2226–47.
- Kumar, V and Rohan Mirchandani. 2012. "Increasing the ROI of social media marketing." *MIT sloan management review* 54 (1):55.
- Lieberman, Andres, Daniel Paravisini, and Vikram Pathania. 2017. "High-Cost Debt and Borrower Reputation: Evidence from the UK." .
- Manski, Charles F. 1993. "Identification of endogenous social effects: The reflection problem." *The review of economic studies* 60 (3):531–542.
- Mas, Alexandre and Enrico Moretti. 2009. "Peers at work." *American Economic Review* 99 (1):112–45.
- McPherson, Miller, Lynn Smith-Lovin, and James M Cook. 2001. "Birds of a feather: Homophily in social networks." *Annual review of sociology* 27 (1):415–444.
- Mikolov, Tomas, Wen-tau Yih, and Geoffrey Zweig. 2013. "Linguistic regularities in continuous space word representations." In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 746–751.
- Mobius, Markus M, Paul Niehaus, and Tanya S Rosenblat. 2005. "Social learning and consumer demand." *Harvard University, mimeograph*. December .
- Moretti, Enrico. 2011. "Social learning and peer effects in consumption: Evidence from movie sales." *The Review of Economic Studies* 78 (1):356–393.
- Mullainathan, Sendhil and Jann Spiess. 2017. "Machine learning: an applied econometric approach." *Journal of Economic Perspectives* 31 (2):87–106.
- Nevo, Aviv. 2001. "Measuring market power in the ready-to-eat cereal industry." *Econometrica* 69 (2):307–342.
- Oster, Emily and Rebecca Thornton. 2012. "Determinants of technology adoption: Peer effects in menstrual cup take-up." *Journal of the European Economic Association* 10 (6):1263–1293.

- Ouimet, Paige and Geoffrey Tate. 2017. "Learning from coworkers: Peer effects on individual investment decisions." Tech. rep., National Bureau of Economic Research.
- Peysakhovich, Alexander and Dean Eckles. 2017. "Learning causal effects from many randomized experiments using regularized instrumental variables." *arXiv preprint arXiv:1701.01140* .
- Piazzesi, Monika, Martin Schneider, and Johannes Stroebel. 2015. "Segmented housing search." Tech. rep., National Bureau of Economic Research.
- Roehm, Michelle L and Alice M Tybout. 2006. "When will a brand scandal spill over, and how should competitors respond?" *Journal of Marketing Research* 43 (3):366–373.
- Sacerdote, Bruce. 2001. "Peer effects with random assignment: Results for Dartmouth roommates." *The Quarterly journal of economics* 116 (2):681–704.
- . 2011. "Peer effects in education: How might they work, how big are they and how much do we know thus far?" In *Handbook of the Economics of Education*, vol. 3. Elsevier, 249–277.
- Sahni, Navdeep S. 2016. "Advertising spillovers: Evidence from online field experiments and implications for returns on advertising." *Journal of Marketing Research* 53 (4):459–478.
- Shapiro, Bradley T. 2018. "Positive Spillovers and Free Riding in Advertising of Prescription Pharmaceuticals: The Case of Antidepressants." *Journal of Political Economy* 126 (1):381–437. URL <https://doi.org/10.1086/695475>.
- Sinkinson, Michael and Amanda Starc. 2018. "Ask Your Doctor? Direct-to-Consumer Advertising of Pharmaceuticals." *Review of Economic Studies* .
- Tucker, Catherine. 2008. "Identifying formal and informal influence in technology adoption with network externalities." *Management Science* 54 (12):2024–2038.
- Valente, Thomas W. 1996. "Network models of the diffusion of innovations." *Computational & Mathematical Organization Theory* 2 (2):163–164.
- Watts, Duncan J, Jonah Peretti, and Michael Frumin. 2007. *Viral marketing for the real world*.
- Young, H Peyton. 2009. "Innovation diffusion in heterogeneous populations: Contagion, social influence, and social learning." *American economic review* 99 (5):1899–1924.

APPENDIX FOR “PEER EFFECTS IN PRODUCT ADOPTION”

Michael Bailey Drew Johnston Theresa Kuchler Johannes Stroebel Arlene Wong

A.1 Random Phone Loss Instrument

In this section, we provide more details about our approach to identifying public posts about random phone loss events. We also provide evidence that random phone loss shocks are not correlated across individuals and their friends.

A.1.1 Random Phone Loss Classification

In order to construct our primary instrument, we need to identify users who have posted publicly about a random phone loss. We take two different approaches to this classification: the first applies regular expression searches, while the second uses machine learning techniques. We find that both classifiers perform well at identifying relevant posts, but that the machine learning-based approach is superior to the regular expression-based classifier in terms of reducing both Type I and Type II errors. As a result, in the paper, we construct the random phone loss instrument using posts identified through the machine learning-based classifier.

Regular Expression Classifier. To build our regular expression-based classifier, we first compiled a list of common phrases (such as “broke my phone” or “phone got stolen”) that were frequently used in Facebook posts concerning lost or broken phones. A complete list of phrases is provided in Table A.1. We then automatically scanned all public Facebook posts by individuals in our sample during the period of our study, flagging those posts that contained at least one of the phrases on our list.

Using this methodology to construct the instruments generates a strong first stage: about 9% of those individuals whose post is flagged end up purchasing a phone in the week of the post. Nevertheless, the classifier identifies a number of false positives (e.g., “So...I dropped my phone in the toilet yesterday...!! Still works tho!!”) while failing to identify a number of more creative descriptions of random phone loss (e.g., “R.I.P phone. You will be missed”). These descriptions are often picked up by our second classification model, described below, which uses natural language processing algorithms.

Machine Learning Classifier. Our machine learning classifier is based on word embeddings, which allow us to translate the unstructured text of the public Facebook posts into features appropriate for machine learning models. Word embeddings are one of the most common tools used in Natural Language Processing (NLP), a sub-field of machine learning that aims to extract insights from data expressed in a human language. They are designed to express a word as a real-valued vector, with the direction and magnitude of each word vector learned from a set of training data. The size of the resulting vectors can vary across implementations, but vectors of 100–1,000 dimensions are commonly used; in our implementation, we use a 200-dimensional embedding. Even though these dimensions are often not easily interpretable, the geometry of the vectors represents semantic and syntactic features of each word, such as tense or quantity (see Mikolov, Yih, and Zweig, 2013). Similar words tend to be represented by similar vectors (as measured by cosine similarity), and linear combinations of

Table A.1: Regular Expressions for Post Classification

Lost Phones			
%lost phone%	%lost iphone%	%lost cell%	%lost my iphone%
%lost iphone%	%lost my phone%	%lost my cell%	
Dropped Phones			
%dropped phone%	%dropped my phone%	%phone dropped%	%cell was dropped%
%dropped iphone%	%dropped my iphone%	%cell dropped%	%phone got dropped%
%dropped cell%	%dropped my cell%	%phone was dropped%	%cell got dropped%
Broken Phones			
%phone%broke%	%broke my iphone%	%brokenphone%	%cellisbroke%
%broke phone%	%broke my cell%	%brokeniphone%	%cell?s broke%
%broke iphone%	%broken cell%	%brokencell%	%brokecell%
%broken phone%	%cell broke%	%cells broke%	%broke my phone%
%broken iphone%	%cellisbroke%	%celljustbroke%	
Destroyed Phones			
%destroyed phone%	%destroyed my cell%	%phone is destroyed%	%cell got destroyed%
%destroyed iphone%	%phone destroyed%	%cell is destroyed%	%destroyed cell%
%cell destroyed%	%phone got destroyed%	%phonedestroyed%	%destroyed my phone%
%phones destroyed%	%cell got destroyed%	%destroyed my iphone%	%cells destroyed%
%phone was destroyed%	%celldestroyed%		
Killed Phones			
%killed phone%	%phoneisdead%	%phone is dead%	%cell just died%
%killed iphone%	%celldead%	%cell is dead%	%phone got killed%
%killed cell%	%cellisdead%	%phone has died%	%killed my phone%
%phone dead%	%cell has died%	%cell got killed%	%killed my iphone%
%cell dead%	%phone died%	%killed my cell%	%phones dead%
%cell died%	%phone was killed%	%phonedead%	%cells dead%
%phone just died%	%cell was killed%		
Smashed Phones			
%smashed phone%	%smashed my iphone%	%phone smashed%	%phone is smashed%
%smashed iphone%	%smashed my cell%	%cell smashed%	%cell is smashed%
%smashed cell%	%phonesmashed%	%phones smashed%	%phone was smashed%
%smashed my phone%	%cellsmashed%	%cells smashed%	%cell was smashed%
Shattered Phones			
%shattered phone%	%shattered my iphone%	%phone shattered%	%phone is shattered%
%shattered iphone%	%shattered my cell%	%cell shattered%	%cell is shattered%
%shattered cell%	%phone shattered%	%phones shattered%	%phone was shattered%
%shattered my phone%	%cell shattered%	%cells shattered%	%cell was shattered%
Busted Phones			
%busted phone%	%busted my cell%	%phone busted%	%phone is busted%
%busted iphone%	%busted my iphone%	%cell busted%	%cell is busted%
%busted cell%	%phone busted%	%phones busted%	%phone was busted%
%busted my phone%	%cellbusted%	%cells busted%	%cell was busted%
Damaged Phones			
%damaged phone%	%damaged my iphone%	%cell damaged%	%phone was damaged%
%damaged iphone%	%damaged my cell%	%phone got damaged%	%cell was damaged%
%damaged my phone%	%phone damaged%	%cell got damaged%	
Stolen Phones			
%stole my phone%	%phone stolen%	%cell got stolen%	%stole my iphone%
%cell stolen%	%phone was stolen%	%stole my cell%	%phone got stolen%
%cell was stolen%			
Not Working Phones			
%phone stopped working%	%phone not working%	%cells not working%	%cell is not working%
%cell not working%	%phone isn?t working%	%cell stopped working%	%phones not working%
%phone is not working%	%cell isn?t working%		
Other			
%phoneless%	%broke%screen%	%shattered%screen%	%contact me here%
%cellless%	%screen%smashed%	%screen%crack%	%no phone%
%smashed%screen%	%crack%screen%	%you can reach me%	%screen%broke%
%screen%shattered%	%contact me on%	%hit me up on here%	

Note: Table shows the regular expressions used to flag posts about random phone loss. % is a wildcard capturing any number of characters (including 0), ? is a wildcard for any single character.

word-embedding vectors contain syntax and semantic meaning. For instance, after converting words to their embeddings, the embedding most similar to $(\vec{King} - \vec{Man} + \vec{Woman})$ is \vec{Queen} .

Several approaches can be taken to learning these vectors from a corpus of text. We chose a skip-gram-based approach, in which a neural network is trained to predict the words surrounding a given term in the corpus. No information is provided to the model about English grammar or syntax—all language features are learned directly from the corpus of text. To train our model, we chose the entire English-language version of Wikipedia as our corpus. Wikipedia is a common corpus in the NLP literature, since it covers a broad range of topics in considerable detail, a feature which helps to train a general-purpose set of embeddings (Bojanowski et al., 2016). We train the embeddings using FastText, a popular open-source library created by researchers at Facebook. After training the word embeddings, we have a model that can transform any word into a 200-dimensional vector. We then concatenate all vectors corresponding to the words in a post, creating a $200 \times N$ matrix that represents the post, where N is the number of words in the post.

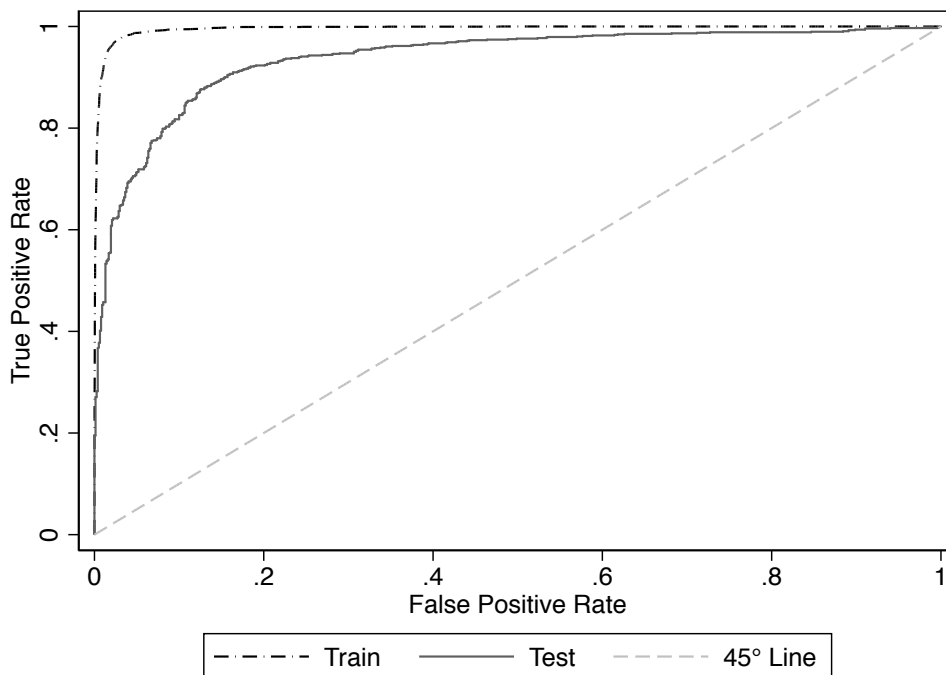
In the next step, we train a convolutional neural network to classify these matrices. Convolutional neural networks (CNNs) are commonly used in natural language processing, as they allow for the creation of very flexible non-linear models that can capture sentence context.¹ This is important for our task, since the ability to distinguish between sentences like “I broke my friend’s phone” and “My friend broke my phone” is crucial. This distinction would have been hard to capture in simpler text classification models that do not respect word order (often called “bag of words” approaches). Convolutional neural networks differ from multi-layer perceptrons (or “vanilla neural networks”) in several ways that are useful for working with text data (Kim, 2014). Specifically, CNNs create convolutional filters that transform the underlying data between traditional layers of the neural network. These filters alter the data to amplify the features that are most relevant in the final classification step. The exact features captured by the filters are determined automatically during the training process. In text data, the filters usually capture and transform multi-word patterns. We employ convolutions of widths 2, 3, 4, 5, and 10 in our work. In general, these convolutions are effective at preserving medium-distance relationships between words, allowing the algorithm to distinguish between phrases like “my phone,” “his phone,” and “a phone.” CNNs can also employ max pooling, which is the selective dropping of data perceived by the neural network to be unimportant. Max pooling normally occurs after a layer of convolutions. This step is important when working with text data, since the dimensions of the input matrix vary between observations. Max pooling allows the model to effectively drop data throughout the process until an appropriately-sized array of features is obtained for the final classification step.

We provided substantial training data to create the final model to identify posts related to random phone loss. To this end, we hand-classified around 8,000 posts picked up by the regular expression model; hand classification identified about 40% of the posts as false positives. We added another 1,000 hand-classified posts that referenced phones in some way but that were not picked up by the regular expression classifier; hand classification revealed that about 20% of these posts concerned a random phone loss, and therefore they represent false negatives for the regular expression classifier.

¹In recent years, CNNs using word embeddings as input have been seen as the state of the art in sentence classification, achieving many of the highest accuracies on common benchmarking tasks (Kim, 2014).

We supplemented these posts with 25,000 posts that had nothing to do with cell phones. The trained model achieved a 0.94 ROC–AUC on 1,800 labeled examples that were set aside as a test set, and which were therefore not used in training the model (see Figure A.1).² To create our primary instrument, we then used the trained model to classify all public posts on Facebook in our sample weeks, identifying those posts that the algorithm detected as being about breaking or losing a phone.

Figure A.1: ROC Curves for Post-Detection CNN



Note: Figure shows the ROC curves for our post-detection CNN. Both the training curve and the test curve are pictured.

Due to the severe class imbalance in our classification problem, some false positives remained with the machine learning classifier. In most cases, the classifier expressed uncertainty about these posts, with estimated probabilities of between 10% and 30% that the posts concerned a random phone loss. Posts to which the algorithm assigned probabilities above 0.3, on the other hand, are almost always true positives. We thus find that we achieve a stronger first stage if we use the neural network’s outputs in tandem with those from the regular expression, in order to give a second check on false positives. In our final model, we therefore set $1(RandomPhoneLoss_{i,t}) = 1$ if the regular expression condition was true and the CNN’s estimated probability was higher than 0.1, or if the regular expression conditions was false and the CNN had an estimated probability above 0.3. This methodology, which is inspired by the notion of ensemble classifiers in machine learning, substantially reduces our false positive rate and

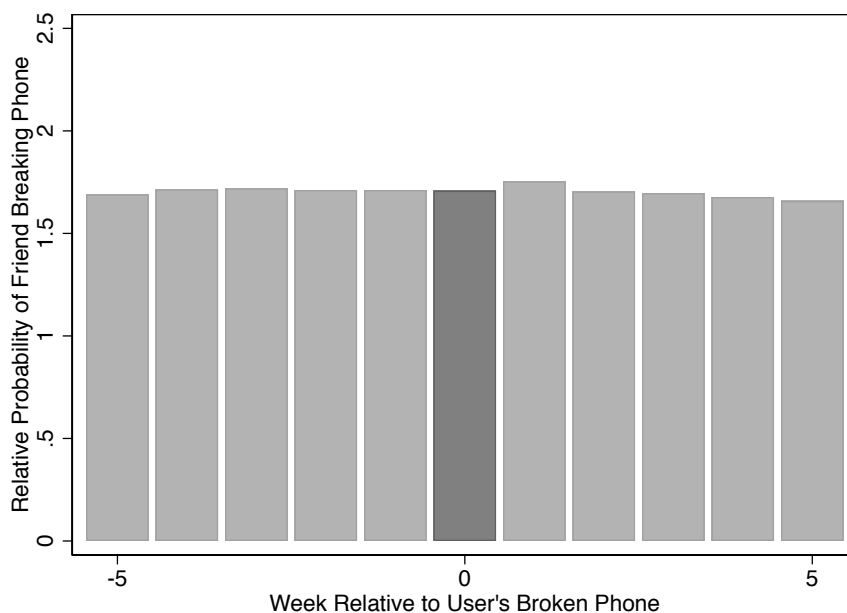
²ROC–AUC is a common metric used in machine learning to evaluate the performance of a predictive model. It can be calculated by plotting a graph of a model’s true positive rate with respect to the false positive rate across all threshold scores and finding the area under the curve of the line formed. Intuitively, it corresponds to the probability a classifier will rank a random positive example above a random negative example (when tasked with distinguishing positive and negative examples). Regardless of the class balance, a score of 0.5 is awarded to a meaningless model (one as good as random guessing), while 1.0 is a perfect score.

generates a stronger first stage than was possible with either of the two classifiers used individually.³

A.1.2 Random Phone Loss Instrument Independence

An important assumption in our empirical analysis is that random phone loss events are not correlated across individuals and their friends. Figure A.2 provides support for this assumption. To construct this graph, we consider two groups: users who post about breaking their phone in week $t = 0$ and users who do not make any such post. For each group, we calculate the average percentage of their friends' posts that are about breaking a phone in weeks $t - 5$ to $t + 5$. We then graph the ratio of the percentages for the first and second group. We see that, even though users who post about breaking their phone tend to have more friends who themselves post about breaking their phone, this feature is constant across time. There does not seem to be any indication that users and their friends tend to disproportionately break their phones at the same time. We control for the fact that different groups tend to post about breaking their phones at different rates by controlling for the number of the user's friends who break their phone in the 12 months prior to our sample. We also control for the average level of $\widehat{ProbBuyRandomPhoneLoss}$ among this group.

Figure A.2: Random Phone Loss Among Friends Relative to User Random Phone Loss



Note: Figure shows the probability that friends of a user who breaks their phone in week 0 post about breaking their own phones in the weeks before and after. The probability is expressed relative to that of friends of users who do not post about a random phone loss.

³The thresholds may appear to be low, but we chose a threshold that balanced the number of the posts found with the conditional probability of switching of the posters. Increasing the thresholds we use would increase the certainty that any individual post is truly a post about a user who breaks their phone, but it would cause the number of potential posts found to be smaller, which negatively impacts the strength of our first stage. We found that using a lower threshold (which permits somewhat more Type II errors in order to reduce the number of Type I errors) gave us the strongest first stage, although raising the threshold somewhat does not change our main findings.

A.2 Propensity Predictions

Throughout this paper, we use several predicted probabilities to construct both our instruments and our controls. In a broad sense, all of these classifiers have a similar goal: we aim to predict users’ phone-purchasing behavior based on information about their demographics, their social networks, and their current phones. The relationships between these features and phone purchasing are complicated and non-linear, so we fit neural networks to predict these propensities using flexible functional forms. We train these models using data from weeks 2016-13 to 2016-15 and 2016-25 to 2016-27 (where weeks are indicated in the format yyyy-ww),⁴ and we then use the models to predict propensities for all observations in the main sample, which runs from 2016-19 to 2016-22. We use five-fold cross-validation to select the hyperparameters for each model. The selected hyperparameters, as well as the resulting in-sample and out-of-sample ROC–AUC scores, are summarized in Table A.2.

Table A.2: Instrument Fitting

	Test Set ROC–AUC	Regression Sample ROC–AUC	Best Layer Sizes
$\widehat{ProbBuyRandomPhoneLoss}_{i,(t,t+1)}$	0.68	0.67	[10]
$\widehat{ProbBuy2y}_{i,(t,t+1)}$	0.64	0.63	[20,5]
$\widehat{ProbUncond}_{i,t}$	0.71	0.70	[10,5]
$\widehat{ProbBuyRandomPhoneLoss}^c_{i,(t,t+1)}$	0.72	0.72	[10,5]
$\widehat{ProbUncond}^c_{i,(t,t+1)}$	0.79	0.78	[20,5]
$\widehat{ProbUncond}^p_{i,(t,t+1)}$	0.80	0.80	[20]

Note: Table shows summary statistics on the predictive power of the classifiers used and the best layer size for the classifier with the best performance on a validation sample. All hyperparameters are determined by five-fold cross-validation. The Test Set ROC–AUC is calculated using a held-out test data set drawn from the same sample as the training data (weeks 2016–13 to 2016–15 and 2016–25 to 2016–27). The Regression Sample ROC–AUC is calculated using data from the main period studied in our regressions (weeks 2016–19 to 2016–22). The scores for the final three groups (which each have several classifications for the different brand or phone model) refer to the averaged one-versus-all ROC–AUC scores for each possible classification.

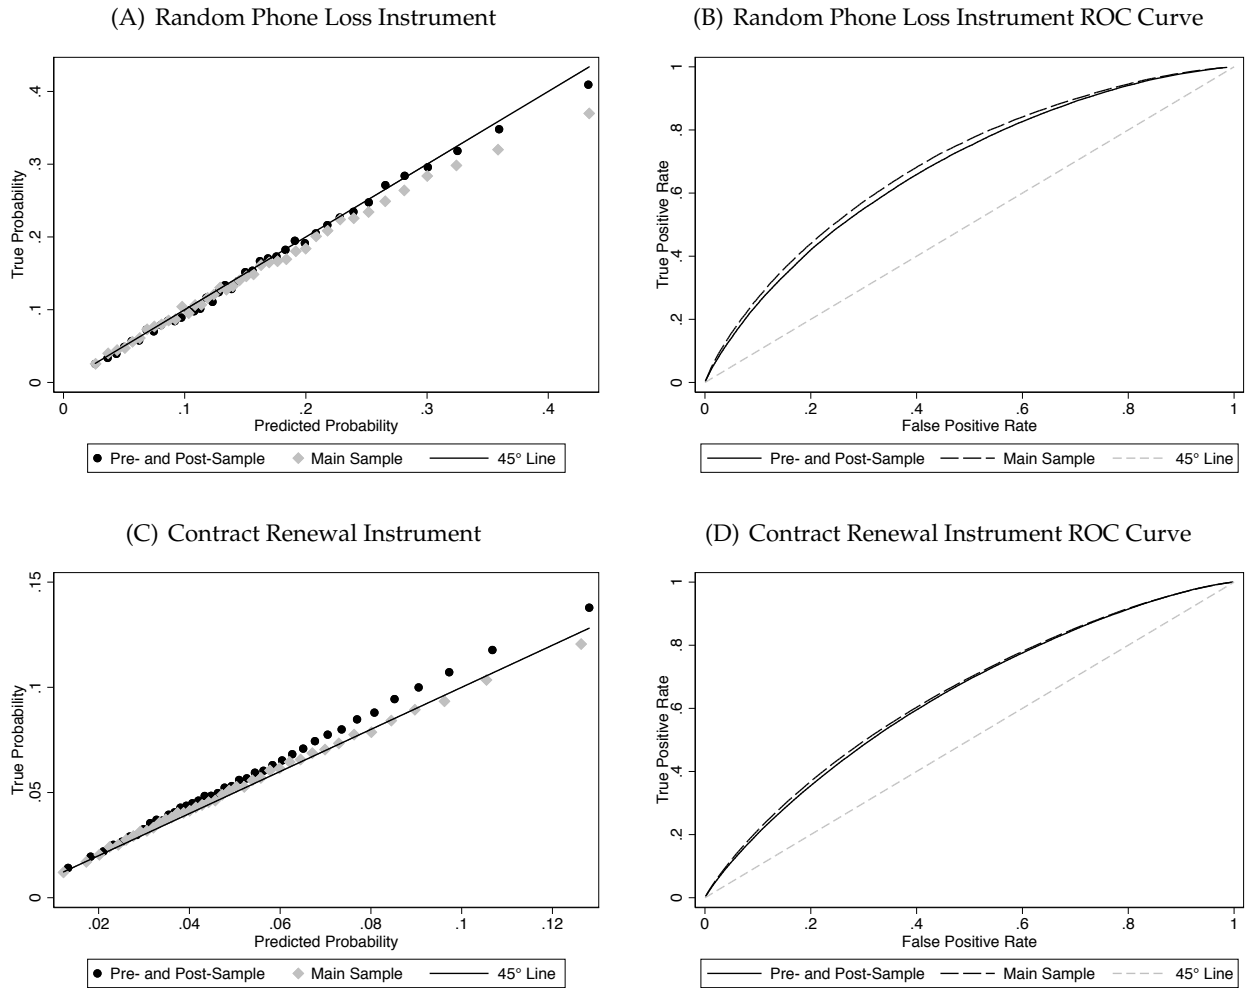
We train both conditional and unconditional models. In the unconditional models, we predict users’ purchasing decisions in weeks t and $t + 1$ on the basis of their characteristics in week t . For all models, the set of observable characteristics used to train the model is as follows: current phone age,⁵ current phone model, carrier, age, user browser, Instagram usage flag, state, education level, friend count, relationship status, activity flags, account age, profile picture flag, number of friendships initiated, gender, and area average income. Since the models are unconditional, we use all user-weeks in our training period as training data. We then predict unconditional probabilities for all users in our main sample, and we use these unconditional probabilities as controls in the regressions in Section 3.

In the conditional models, we aim to predict a user’s probability of purchasing a phone in weeks

⁴Training on held-out data is important in this case, as training using in-sample data would run the risk of overfitting the model, which would bias the IV estimates towards the OLS coefficients.

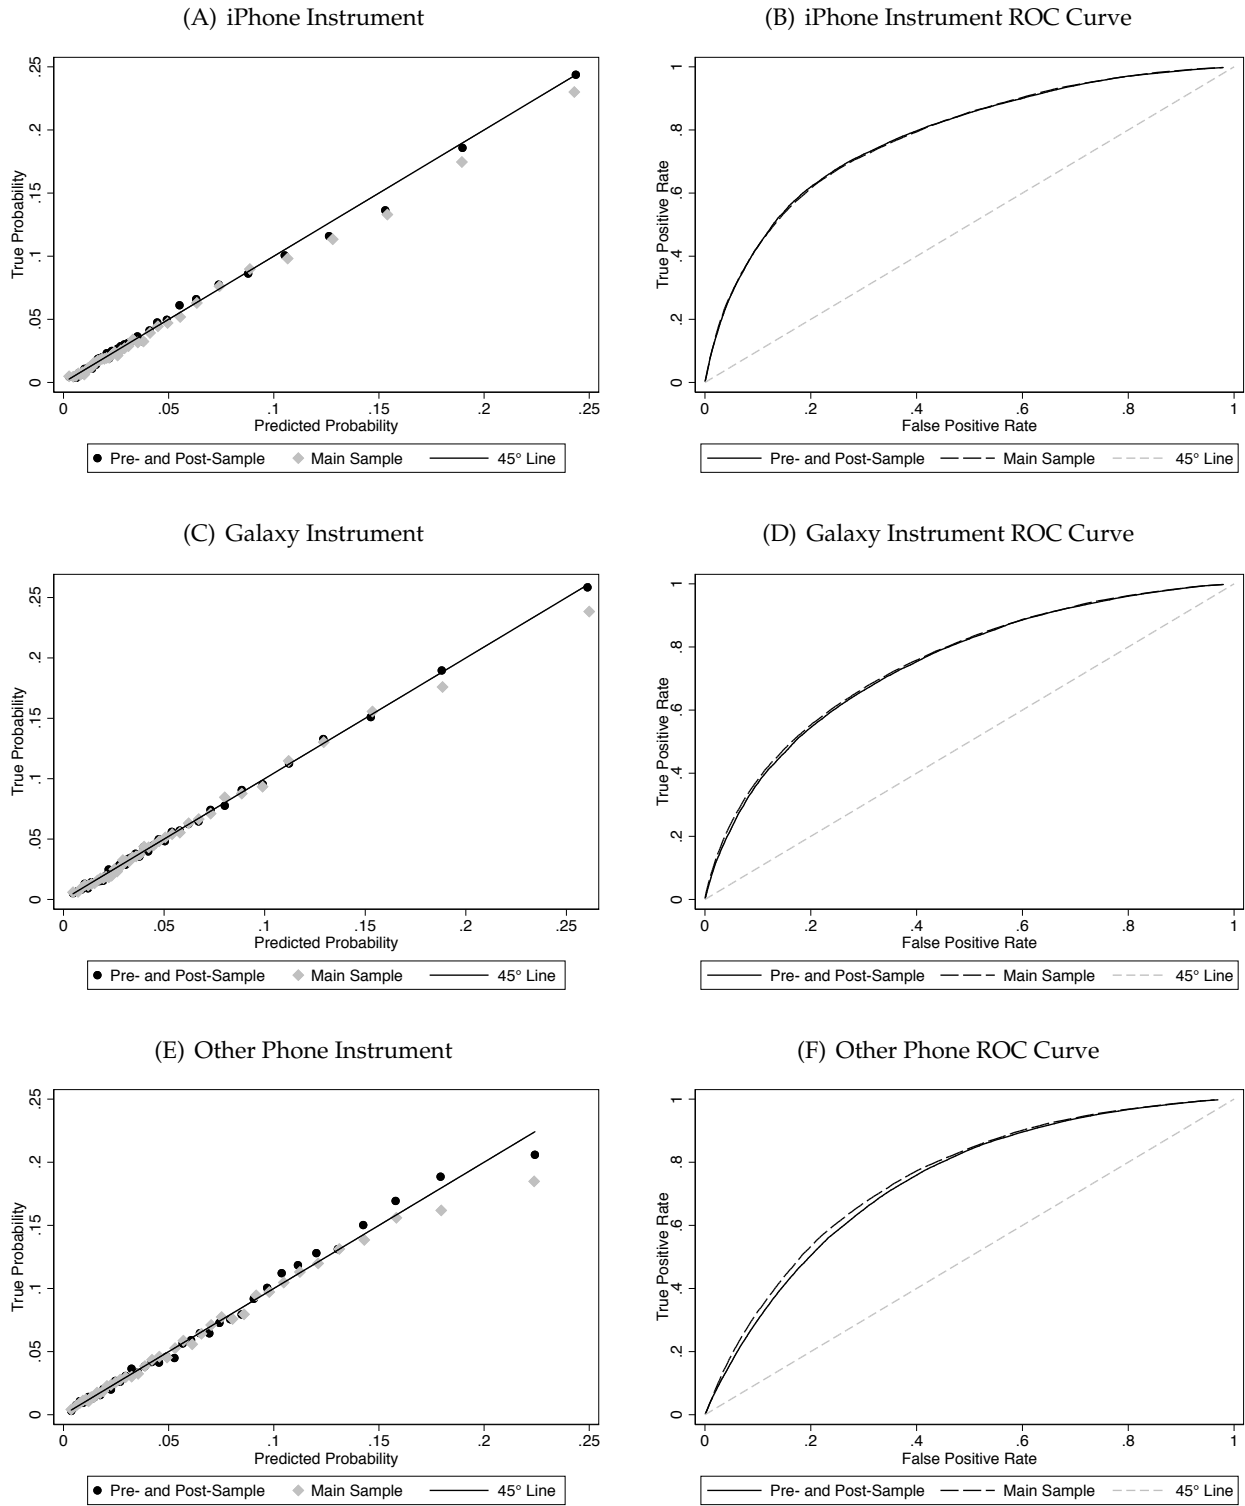
⁵We do not include this characteristic in our models studying switching probability at the contract renewal threshold, as all users in the training set for this classifier have their phone age in a narrow range. This feature makes predicting the conditional probability for users whose phone age is outside this range difficult.

Figure A.3: Baseline Instruments



Note: Panels A and C in the left column show bincscatter plots of the fit of our two instruments. Panels C and D in the right column present Receiver Operating Characteristic curves for each of the instruments. All plots only include users for whom $\mathbb{1}(Instrument_{i,t}) = 1$. The regression results using these instruments are shown in Table 2. The “Pre- and Post Sample” is a held-out set of observations from weeks 2016–13 to 2016–15 and 2016–25 to 2016–27, the same weeks that were used to train the data. The “Main Sample” is a random sample of 10 million observations from the period 2016–18 to 2016–22, the period used to construct our main panel.

Figure A.4: Brand Instruments



Note: The left column shows binscatter plots of the fit of our probabilistic instruments for brand-level phone purchases. The right column presents Receiver Operating Characteristic curves for each of these instruments. All plots only include users for whom $\mathbb{1}(BrokenPhone_{j,t}) = 1$. The regression results using these three instruments are shown in Table 3. The “Pre- and Post Sample” is a held-out set of observations from weeks 2016–13 to 2016–15 and 2016–25 to 2016–27, the same weeks that were used to train the data. The “Main Sample” is a random sample of 10 million observations from the period 2016–18 to 2016–22, the period used to construct our main panel.

t and $t + 1$, conditional on some behavior or trait in week t , such as posting about a random phone loss or having a phone at the contract renewal threshold. We fit one such model for each of the two instruments, and we predict each conditional probability for all user-weeks in our main sample. In addition to the summary statistics on the model fit provided in Table A.2, Figure A.3 provides additional information on the performance of the various classifiers for our three baseline instruments. The left column shows binscatter plots of the predicted probability against the realized probability for both the training data and the regression data set. Reassuringly, the training data line up on the 45-degree line, and the regression sample data, which were not used to train the model, also align closely. This finding suggests that our models are stable over time. The horizontal axis shows the range of predicted probabilities. For example, we find the probabilities for purchasing in weeks t and $t + 1$ after posting about a random phone loss in week t range to be between 5% and 50%. This result highlights the value of using our instrumental variables approach—rather than just using the number of friends with a random phone loss—as instruments for the number of friends buying a phone.

To perform the regressions in Section 4, we additionally fit models that predict acquisitions of particular phone types. The same features used to generate the general switching predictions are highly predictive of which phone a user buys, and our models here have strong predictive power. We create two granularities of classifiers. In the first case, we train a classifier to predict a user’s probability of buying any of three mutually exclusive and exhaustive categories of cell phones, iPhones, Galaxies, and other phones, as well as their probability of not buying any phone. We use these predicted probabilities as described in Section 4.1. In the latter case, described in Section 4.2, we predict the user’s probability of purchasing each of the 20 most commonly-purchased phones, three residual categories (other iPhone, other Galaxy, other), and no phone at all. For both granularities, we train unconditional models for use as controls in our regressions: $\widehat{ProbUncond}_{i,t}^c$ for the brand-level granularity, and $\widehat{ProbUncond}_{i,t}^p$ for the model-level granularity. At the brand level, we also train a classifier conditional on a user breaking or losing their phone, $\widehat{ProbLost}_{i,t}^c$; the classifiers are used to construct our instruments in Section 4.1. Figure A.4 explores the performance of these predictors. We do not train conditional models at the model level, since we find that the predictions within a brand tend to be highly collinear, making first stage estimation complicated (see Section 4.2 for further discussion). Tables A.2 and A.3 show the performance of the brand-level and model-level propensity predictions.

Table A.3: Phone-Specific ROC–AUC

	Unconditional		Conditional on Losing Phone	
	Test Set	Regression Sample	Test Set	Regression Sample
	ROC–AUC	ROC–AUC	ROC–AUC	ROC–AUC
Brand	0.755	0.752	0.720	0.716
iPhone	0.734	0.734	0.757	0.758
Galaxy	0.772	0.769	0.729	0.726
Other	0.808	0.799	0.711	0.707
No Purchase	0.708	0.704	0.684	0.673
Specific	0.786	0.782		
iPhone 6S	0.782	0.785		
iPhone 6S Plus	0.711	0.718		
iPhone SE	0.773	0.781		
iPhone 6	0.738	0.733		
iPhone 5S	0.720	0.716		
iPhone 6 Plus	0.680	0.688		
Other iPhone	0.706	0.689		
Galaxy S7	0.820	0.816		
Galaxy S7 Edge	0.790	0.786		
Galaxy Core Prime	0.797	0.793		
Galaxy Note 5	0.808	0.809		
Galaxy J7	0.851	0.859		
Galaxy Grand Prime	0.793	0.795		
Galaxy S5	0.758	0.741		
Galaxy S6	0.763	0.755		
Other Galaxy	0.767	0.758		
Tribute 5	0.857	0.847		
K10	0.844	0.846		
G5	0.830	0.810		
G Stylo	0.851	0.844		
Desire 626s	0.852	0.839		
One Touch	0.881	0.875		
Other	0.800	0.793		
No Purchase	0.702	0.700		

Note: Table shows summary statistics on the classifiers used to predict the phone type a user purchases. Since ROC–AUC is a score for binary classification problems, the multi-class ROC–AUC scores (bolded) represent the averaged one-versus-all ROC–AUC scores of each individual classification problem. We do not display conditional ROC–AUC scores for specific phones, as we report only OLS versions of these regressions due to the high correlation between within-brand scores as discussed further in Section 4.2.

A.3 Heterogeneity Models

In Section 3.3, we estimate heterogeneities in influence and susceptibility to influence across individuals and relationships. In this Appendix, we provide the regression specifications we use in these analyses.

Heterogeneity by Relationship and Friend Characteristics. We first analyze heterogeneities in influence according to relationship characteristics and friend characteristics. In each regression, we consider a mutually-exclusive and exhaustive group of characteristics G . For relationship characteristics, one group of characteristics corresponds to family and non-family relationships. For friend characteristics, one group of characteristics corresponds to friend ages, with three conditions $g \in G$ capturing friends aged 18–25, 26–40, and 40+. We use these conditions to create new instruments and endogenous variables for each $g \in G$:

$$Instrument_{i,t}^{Lose,g} = \sum_{j \in Fr(i)} \widehat{ProbBuyRandomPhoneLoss}_{j,t} \cdot \mathbb{1}(LostPhone_{j,t}) \cdot \mathbb{1}(Condition_{j,t}^g) \quad (A.1)$$

$$FriendsBuyPhone_{i,(t,t+1)}^g = \sum_{j \in Fr(i)} \mathbb{1}(BuysPhone)_{j,(t,t+1)} \cdot \mathbb{1}(Condition_{j,t}^g) \quad (A.2)$$

We also create two new sets of controls. The first set counts the number of friends of user i who are members of each group g , while the second controls for the average conditional probability among the user's friends in each group:

$$Friends_{i,t}^g = \sum_{j \in Fr(i)} \mathbb{1}(Condition_{j,t}^g) \quad (A.3)$$

$$FriendAvgProbLost_{i,t}^g = \frac{\sum_{j \in Fr(i)} \widehat{ProbBuyRandomPhoneLoss}_{j,t}^g}{Friends_{i,t}^g} \quad (A.4)$$

Using these new variables, we estimate one first stage per condition using the following first and second stages:

$$FriendsBuyPhone_{i,(t-1,t)}^g = \sum_{g \in G} \delta_g Instrument_{i,t-1}^{Lose,g} + \sum_{g \in G} \psi_g Friends_{i,t-1}^g + \sum_{g \in G} \theta_g FriendAvgProbLost_{i,t-1}^g + \omega X_{i,t} + e_{i,t} \quad (A.5)$$

$$\mathbb{1}(BuysPhone_{i,t}) = \sum_{g \in G} \beta_g FriendsBuyPhone_{i,(t-1,t)}^g + \sum_{g \in G} \Psi_g Friends_{i,t-1}^g + \sum_{g \in G} \Theta_g FriendAvgProbLost_{i,t-1}^g + \gamma X_{i,t} + \epsilon_{i,t} \quad (A.6)$$

In each case, we include the same fixed effects and controls outlined in column 2 of Table 2.

Heterogeneity by User Characteristics. In Section 3.3, we also analyze how susceptibility to influence varies according to user characteristics. We employ a similar approach to that outlined above to understand these heterogeneities:

$$Instrument_{i,t}^{Lose,g} = \mathbb{1}(Condition_{i,t}^g) \cdot \sum_{j \in Fr(i)} \widehat{ProbBuyRandomPhoneLoss}_{j,t} \cdot \mathbb{1}(LostPhone_{j,t}) \quad (A.7)$$

$$FriendsBuyPhone_{i,(t,t+1)}^g = \mathbb{1}(Condition_{i,t}^g) \cdot \sum_{j \in Fr(i)} \mathbb{1}(BuysPhone)_{j,(t,t+1)} \quad (A.8)$$

Our regressions are similar to those outlined above, but when studying user characteristics, we do not need to include separate controls indicating whether user i meets condition g , because each of these conditions is already included as a fixed effect at a finer level of granularity. Our first and second stages are

$$FriendsBuyPhone_{i,(t-1,t)}^g = \sum_{g \in G} (\delta_g Instrument_{i,t-1}^{Lose,g}) + \omega X_{i,t} + e_{i,t} \quad (A.9)$$

$$\mathbb{1}(BuysPhone_{i,t}) = \sum_{g \in G} (\beta_g Friends\widehat{BuyPhone}_{i,(t-1,t)}^g) + \gamma X_{i,t} + \epsilon_{i,t} \quad (A.10)$$

In each case, we include the same fixed effects and controls outlined in column 2 of Table 2.

Heterogeneity by Pairwise Characteristics In Section 3.3, we also study variation in peer effects according to user and friend characteristics simultaneously. We consider the same set G of characteristics for both the users and their friends. We construct new instruments and endogenous variables as follows:

$$Instrument_{i,t}^{Lose,g_1,g_2} = \mathbb{1}(Condition_{i,t+1}^{g_1}) \cdot \sum_{j \in Fr(i)} \mathbb{1}(ProbBuyRandomPhoneLoss)_{j,t} \cdot \mathbb{1}(Condition_{j,t}^{g_2}) \quad (A.11)$$

$$FriendsBuyPhone_{i,(t,t+1)}^{g_1,g_2} = \mathbb{1}(Condition_{i,t+1}^{g_1}) \cdot \sum_{j \in Fr(i)} (\mathbb{1}(BuysPhone)_{j,(t,t+1)} \cdot \mathbb{1}(Condition_{j,t}^{g_2})) \quad (A.12)$$

We use these variables to construct first stages that allow us to separately gauge the influence of friends of each type on users of each type. The regressions for the first and second stages, respectively, are

$$FriendsBuyPhone_{i,(t-1,t)}^{g_1,g_2} = \sum_{g_1 \in G} \sum_{g_2 \in G} (\delta_{g_1,g_2} Instrument_{i,t-1}^{Lose,g_1,g_2}) + \sum_{g \in G} \psi_g Friends_{i,t-1}^g + \sum_{g \in G} \theta_g Friend\widehat{AvgProbLost}_{i,t-1}^g + \omega X_{i,t} + e_{i,t} \quad (A.13)$$

$$\mathbb{1}(BuysPhone_{i,t}) = \sum_{g_2 \in G} \sum_{g_1 \in G} (\beta_{g_1,g_2} Friends\widehat{BuyPhone}_{i,(t-1,t)}^{g_1,g_2}) + \sum_{g \in G} \Psi_g Friends_{i,t-1}^g + \sum_{g \in G} \Theta_g Friend\widehat{AvgProbLost}_{i,t-1}^g + \gamma X_{i,t} + \epsilon_{i,t} \quad (A.14)$$

In each case, we the same fixed effects and controls outlined in column 2 of Table 2. For each individual, we additionally control for the number of friends in each group.