

PEER EFFECTS IN PRODUCT ADOPTION*

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

Abstract

We use de-identified data from Facebook to study the nature of peer effects in the market for cell phones. To identify peer effects, we exploit variation in friends' new phone acquisitions resulting from random phone losses. A new phone purchase by a friend has a large and persistent effect on an individual's own demand for phones of the same brand. While peer effects increase the overall demand for phones, a friend's purchase of a particular phone brand can reduce an individual's own demand for phones from competing brands, in particular if they are running on a different operating system.

JEL Codes: L1, L2, M3, D4

Keywords: Peer Effects, Demand Spillovers, Social Learning

*This version: January 18, 2021. For their helpful comments, we thank the editors, Camille Landais and Seema Jayachandran, several anonymous referees, as well as Tim Armstrong, John Campbell, Dean Eckles, Liran Einav, Matt Gentzkow, Ed Glaeser, Paul Goldsmith-Pinkam, Matthew Jackson, 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, Frankfurt School, LMU Munich, NYU, Rochester, Stanford, Stockholm University, UNC Chapel Hill, University of Bonn, University of British Columbia, University of Mannheim, UT Austin, Wharton, Yale, the Empirical Macro Workshop in Las Vegas, and the NBER Summer Institute. 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: mcbailey@fb.com

[‡]Harvard University. Email: drewmjohnston1@gmail.com

[§]New York University, Stern School of Business, NBER, and CEPR. 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

Peer effects in consumption are pervasive. For example, an individual's choice of which car to purchase is likely influenced by the recent car purchasing decisions of her friends. Such peer effects have important implications for firms and policy makers. For instance, in the presence of peer effects, the elasticity of aggregate demand may be larger than the elasticity of individual demand, since any direct incremental sales in response to a price reduction may lead to further extra sales through peer effects. Similarly, from a macro perspective, such peer effects in consumption suggest that the effects of stimulus policies on aggregate demand are larger than those estimated from directly-affected individuals.

Despite the economic importance of peer effects in consumption and product adoption decisions, there is limited evidence on their exact nature and the resulting implications. For example, peer effects may lead someone to buy a new phone when her 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 such peer effects for the competitive dynamics between firms also depend on whether any changes in demand are restricted to the brand 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, heterogeneous, and long-lasting, and that they generate substantial incremental demand. Positive peer effects are largest for the brand purchased by the peer, but the size of the peer effect on same-brand demand often exceeds the effect on total phone demand. This finding suggests that some incremental same-brand purchases come at the expense of purchases from competing brands, in particular those on different operating systems.

We work with de-identified data from Facebook, the world's largest online social networking site. At the end of our sample period in May 2016, Facebook had around 226 million active users in the U.S. and Canada (Facebook, 2016). 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, allowing us to identify the timing of new phone acquisitions.

We use this data 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 purchases. 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.

First, we use the number of friends who break or lose their phones in a given week to instrument for how many friends purchase a new phone in that week. The identifying assumption is that the number of friends who 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 approach 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 buy 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 up for a contract renewal, which is often aligned with an upgrade to a new device.

We improve the power of these instruments by exploiting variation not only in *how many* friends experience the conditionally random event, but also in *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, that older individuals are more likely to buy a new phone immediately after breaking their old device. The eventual instrument, then, 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. This research design allows us to control, for example, for the average age in a person’s friendship network, and only identify off variation in whether it is the person’s old or young friends who happen to break their phones in a given 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 0.040 and 0.022 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. We argue that much of the communication between friends about the new phone purchase that drives the observed peer effect occurs off the Facebook platform, and to a substantial extent through real-world interactions. Consistent with this interpretation, we show that peer effects from geographically proximate friends are larger than peer effects from friends who live further away.

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. Quantitatively, having one extra friend purchase a new phone increases an individual’s own probability of purchasing a new phone over the next 4 months by 0.6 percentage points, relative to a baseline probability of buying a new phone over this horizon of about 14.6%.

In the next step, we explore heterogeneities in peer effects along characteristics of potential influencers and potentially influenced individuals. We focus on heterogeneities in the local average treatment effects of the random phone loss instrument, which has the most power in the baseline specification, but find similar patterns of heterogeneity in the corresponding OLS estimates. We observe that close friends on Facebook exert a larger influence on one another than friends with weaker tie strength. We also find large heterogeneities in the peer effects exerted by different demographic groups, but lit-

the variation in individual susceptibility to influence along the same demographic characteristics. For example, less-educated individuals have the largest effects on their friends' purchasing behaviors, but these individuals are no more likely to be influenced by phone purchases of their friends. These heterogeneities in peer influence have important implications for understanding the effectiveness of seed marketing campaigns, which target a small set of early adopters who can generate follow-on demand through peer effects. 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.

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, and "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 who 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 happen to lose their phones are those who are likely to replace it with a new iPhone or those who are likely to purchase a new Samsung Galaxy.

There are three key take-aways from the cross-brand analysis. First, for all three brand categories, positive peer effects are largest for phones in the same category as that purchased by the peer. Second, 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.¹ The third main take-away relates to across-brand demand spillovers. Specifically, we find that when a friend buys a new phone, this purchase increases a person's own propensity of buying a phone from competing brands on the same operating system, while reducing their propensity of buying a phone from competing brands on different operating systems. In other words, while some of the observed positive same-brand peer effects arise by generating entirely new demand, others come from pulling demand away from rival firms with competing operating systems. Importantly, these demand spillovers across operating systems could have easily been positive. For example, a user who buys a Galaxy might have caused her friends to desire a new phone—of any type, including iPhones—through a "keeping up" effect. The observed across-brand demand spillovers are thus again consistent with an important social learning component: when your friends use a certain operating system, you are more likely to learn about that system. This would increase your demand for all phones using that operating system (even those produced by a different manufacturer), in part at the expense of phones using competing operating systems.²

¹Our results do not allow us to rule out that "keeping up" effects (which are likely to be larger for more expensive brands) also contribute to the observed peer effects; instead, our findings suggest that such effects cannot be the entire story.

²In addition to social learning, network externalities provide a second mechanism that might explain some of the patterns

The observed across-brand demand spillovers highlight 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 may 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., 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 negative across-brand demand spillovers 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.

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., Goolsbee and Klenow, 2002; Mobius, Niehaus, and Rosenblat, 2005; Kuhn et al., 2011; Moretti, 2011; Aral and Walker, 2012; Gilchrist and Sands, 2016; De Giorgi, Frederiksen, and Pistaferri, 2016; Han, Hirshleifer, and Walden, 2016) as well as a variety of household financial decisions (e.g., Duflo and Saez, 2003; Hong, Kubik, and Stein, 2004; Bursztyrn et al., 2014; Beshears et al., 2015; Ouimet and Tate, 2017; Kuchler and Stroebl, 2020), housing market decisions (e.g., Bailey et al., 2019, 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), program participation (Dahl, Løken, and Mogstad, 2014), labor market outcomes (Mas and Moretti, 2009), mutual fund investments (Kuchler et al., 2020), international trade flows (Bailey et al., 2020b), and the spread of and response to COVID-19 (Bailey et al., 2020c; Kuchler, Russel, and Stroebl, 2020). 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; Conley and Udry, 2010; Oster and Thornton, 2012; Kremer and Miguel, 2007; Björkegren, 2018). 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). 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 important 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 influential individuals, as well as the correlation of peer influence with price sensitivity.

of across-brand demand spillovers. Such network externalities would arise if having more friends use a certain operating system would increase a user's own value of using that same operating system. In the context of cell phones, network externalities may primarily come from the use of the FaceTime video messaging app, which is only available on the iOS operating system. However, while we cannot rule out that such network externalities play some role, such externalities cannot explain a number of the patterns we document in this paper, all of which would naturally follow from a social learning story (e.g., the fact that peer effects are largest for the same model rather than equally spread across all phones of the same operating system, or the fact that we find peer effects to decline in time since model release). As a result, our findings are only consistent with a story in which there is at least a substantial social learning component to peer effects (in addition to potential other components coming from "keeping up" desires or network externalities).

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. 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, 2020a).

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.³ 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.⁴ 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 Labor Day), which could confound our estimates. We are left with about 329 million user-weeks as our baseline estimation sample.

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 389 days old, while the median user has a phone that is just over ten months old. The 10th–90th-percentile range of phone age is between 63 days and 777 days. About 0.93% of all users acquire a new phone in a given week. The average user has 323 friends in the sample as well as about 3 new phone purchases among friends in a given week.

³Facebook reports in its July 26, 2018, 10-Q filing: “Substantially all of our daily and monthly active users [...] access Facebook on mobile devices.”

⁴The 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. The overwhelming majority of switches that we detect are to phones released no more than nine months prior to the start of our sample, suggesting they are new purchases rather than hand-downs from friends and family.

Table 1: Summary Statistics

	Mean	Standard Deviation	P10	P25	P50	P75	P90
User Characteristics							
Age (Years)	35.3	12.1	21	25	33	44	53
Male	0.58	0.5	0	0	1	1	1
Phone Age (Days)	388.8	327.3	63	152	317	544	777
Buys Phone (%)	0.93	9.59	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	322.4	202.8	124	165	258	424	631
Friends with Phone Purchases	3.00	2.87	0	1	2	4	7
Friends with Public Statuses	59.5	53.8	17	26	43	73	120
Friends Posting about Breaking/Losing Phone	0.26	0.64	0	0	0	0	1
Friends at Phone Age Threshold	1.83	1.84	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 329 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 behaviors, 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. We next discuss both of these instruments in more detail.

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. As a result, an individual who has more friends randomly losing their phones in a given week is likely to have more friends buying a new phone in that week. Provided that a random phone loss of a friend only influences the probability that a user herself purchases a new phone through peer effects from any replacement purchase by the friend, the number of friends who experience a random phone loss can then be used to instrument for the number of friends who purchase new 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.

Figure 1: Sample Posts About Randomly-Lost Phones



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.⁵ Specifically, we use two tools from the natural language processing literature: word embeddings and convolutional neural networks. We will provide a brief overview of these tools here; a longer explanation of our methodology is available in Appendix A.1.

Our approach relies on word embeddings, which are low-dimensional vectors that provide a geometric representation of the meaning of the corresponding word. Words with similar meanings will be represented by similar vectors, and the spatial relationships between vectors will capture complex

⁵We 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 in our sample has about 322 friends in total, only about 60 of those friends have set their statuses as public.

relationships between the corresponding words (see Mikolov, Yih, and Zweig, 2013, for details). For instance, after converting words to their embeddings, the embedding most similar to $(\overrightarrow{King} - \overrightarrow{Man} + \overrightarrow{Woman})$ is $\overrightarrow{Queen}$. In our application, we use 200-dimensional word embeddings that were trained using all articles on the English edition of Wikipedia. Using these vectors, we can represent each public Facebook post as a matrix, consisting of the stacked vectors of its constituent words.

After generating this numerical representation of each public post, we next use a convolutional neural network (CNN) to determine which posts describes a user breaking or losing her phone. CNNs were originally developed for applications in computer vision, and they expand upon traditional neural networks by transforming the underlying data to make use of its spatial configuration. In the case of image data, CNNs account for relationships between nearby areas of the image; in natural language applications, CNNs make use of the order of words within a passage. This allows us to distinguish between sentences such as *“I broke my phone when I was with my friend John”* and *“I just saw my friend John break his phone”*. We train the CNN on a large sample of manually-classified posts using 10-fold cross-validation, and then use it to classify all public posts in our sample. The resulting model performs quite well on unseen posts, identifying many idiosyncratic examples such as *“R.I.P phone. You will be missed”* that would be difficult to capture with regular expression searches.⁶ Appendix A.1 includes further details on the training process and the model’s performance. In total, we identify around 65,000 public posts about broken or lost phones per week. Table 1 shows that, in a given week, the average person has 0.26 friends who publicly post about breaking or losing their phones.⁷

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 substantial increase in the probability of acquiring a new phone. Specifically, about 10% of individuals with a post identified by our classifier get a new phone in the week of posting about losing their 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 ob-

⁶It is likely that the CNN identified this particular post after observing hand-classified posts such as *“My phone is dead”* in the training sample, combined with the fact that *“dead”* and *“R.I.P”* occupy similar positions in the embedding space.


⁷We have also implemented a model using a regular expression-based classifier, which produced an instrument that had less power but found largely similar results as our baseline analysis. This simpler classifier is used to reinforce our main model in an approach inspired by ensemble classifiers. See the discussion in Appendix A.1.

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 a given week, splitting users by their posting behavior in week 0. The line *Random Phone Loss Post* (green triangles) shows the behavior of users who have a public post in week 0 that relates to a random phone loss. The line *Other Post* (orange circles) 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* (blue squares) 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).

serve a new phone purchase for every individual whom we identify as having posted about a random phone loss. There are several reasons for this result. First, 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.⁸ 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 behavior 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.⁹

While posting about breaking or losing one’s phone leads to a sizable 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.

⁸Properly 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 associated decrease in the number of true positives thus weakened our instrument.

⁹Additionally, 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.

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 potential 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 the large space of possible instruments.

$$ProbBuyRandomPhoneLoss_{i,t} = Prob(\mathbb{1}(BuysPhone_{i,t})|X_{i,t}, \mathbb{1}(RandomPhoneLoss_{i,t}) = 1). \quad (1)$$

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 $ProbBuyRandomPhoneLoss_{i,t}$. 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:


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

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 $ProbBuyRandomPhoneLoss_{j,t}$ 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

¹⁰We use the following characteristics as features when training our neural networks: current phone age, current phone model, carrier, user age, user gender, user browser, Instagram usage flag, user education level, U.S. state, friend count, activity flags, account age, profile picture flag, number of friendships initiated, and area-level average income, where we measure location at the zip code level for location-services users, and at the county level for all other users.

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

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, $ProbBuyUncond_{i,t}$, 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 6, with the exception of $ProbBuyRandomPhoneLoss_{j,t}$, the horizontal axis variable. Panels C and D in the bottom row show the analogous relationships for the contract renewal instrument.

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, $ProbBuyUncond_{i,t}$, 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 top row, we explore the random phone loss instrument (equation 2). Panel A shows unconditional relationships. We find that $Instrument_{i,t}^{Lose}$ is correlated

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.

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—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 with unobservable user characteristics, either.

It is important to point out that the set of compliers in a specification using $Instrument_{i,t}^{Lose}$ to instrument for the total number of phone purchases by friends is likely different to the set of compliers when using the total number of friends who break their phones. To the extent that these compliers differ in the strength of the peer effects they exert, the two approach may therefore estimate different local average treatment effects (LATEs), though it is unclear whether either one of these LATEs would be preferable in terms of being more representative of a population average treatment effect.

2.2 Phone Age Instrument

Our second instrument is based on the observation 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 is generally increasing in phone age, but it spikes when phones cross the two-year age threshold (the dark grey area). Panel B, which shows this probability separately by carrier, 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:

$$ProbBuy2y_{i,t} = Prob(\mathbb{1}(BuysPhone_{i,t}) | X_{i,t}, \mathbb{1}(Phone2yOld_{i,t}) = 1), \quad (3)$$

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 included in $X_{i,t}$ also influence this conditional probability. We then instrument for the number of friends who get a new phone with the sum of $ProbBuy2y_{j,t}$ across all friends who are at the two-year phone age threshold in a given week:

$$Instrument_{i,t}^{2y} = \sum_{j \in Fr(i)} \mathbb{1}(Phone2yOld_{j,t}) \cdot ProbBuy2y_{j,t}. \quad (4)$$

Figure 4: Probability of New Phone by Phone Age



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.

Since individuals who have more friends with older phones are plausibly different from individuals with friends who have younger phones, 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 $ProbBuy2y_{j,t}$ among those people, in addition to the average value of $ProbBuy2y_{j,t}$ 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, $ProbBuyUncond_{i,t}$.

2.3 Empirical Specification and Inference

Using these instruments, we estimate instrumental variables (IV) regressions to measure peer effects in the cell phone market. The first and second stages of the IV regression, respectively, are:

$$FriendsBuyPhone_{i,(t-1,t)} = \delta Instrument_{i,t-1} + \omega X_{i,t} + e_{i,t} \quad (5)$$

$$\mathbb{1}(BuysPhone_{i,t}) = \beta \widehat{FriendsBuyPhone}_{i,(t-1,t)} + \gamma X_{i,t} + \epsilon_{i,t}. \quad (6)$$

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 buckets $\times$ 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, $ProbBuyUncond_{i,t}$. In the Appendix, we show that our baseline results are robust to different specifications of the controls and fixed effects.

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 that 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, $FriendsBuyPhone_{i,(t-1,t)}$:

$$FriendsBuyPhone_{i,(t-1,t)} = \sum_{j \in Fr(i)} \mathbb{1}(BuysPhone)_{j,t-1} + \sum_{j \in Fr(i)} \mathbb{1}(BuysPhone)_{j,t}. \quad (7)$$

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

Inference. In any setting where peer effects might be important (whether or not they are the focus of the analysis), these peer effects can introduce a correlation in the error terms across individuals. Such a correlation would invalidate the independence assumptions used to derive the asymptotic properties of standard estimators. In a world with non-overlapping network communities, one can account for this possible across-observation dependence due to peer effects by clustering standard errors at the level of the community. For complete networks like the one we are studying, statistical inference remains a relatively open area of research, and our vast sample size limits our ability to use the social graph to fully model the structure of the variance-covariance matrix (similar issues arise in a literature that explores the use of cluster-robust estimators when working with spatially dependent data, see Bester, Conley, and Hansen, 2011). We therefore follow a number of recent papers to explore the robustness of our statistical inference to various approaches of constructing standard errors. In particular, Eckles, Kizilcec, and Bakshy (2016) and Zacchia (2020) propose to partition the social graph into a number of communities with limited cross-community dependence, and to then cluster the standard errors at the community level.¹³ Even though the presence of some across-cluster friendship links implies that there remains the potential for across-cluster correlation in the error terms, this clustering approach substantially reduces potential biases in standard errors from such dependencies. Appendix A.3 shows that our standard errors are essentially unaffected when moving from heteroskedasticity-robust standard errors to community-robust standard errors. This suggests that in our setting, statistical inference is not substantially affected by residual across-individual dependencies in error terms.

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. We begin by presenting the baseline estimates of peer effects. Section 3.1 then explores the

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.

¹³Operationally, we start with a data set that uses a distributed variant of the Kernighan-Lin algorithm to divide the global Facebook social graph into about 21,000 distinct communities. Individuals in our sample are assigned to their communities created by this graph. The 0.2% of our sample assigned to communities with fewer than 100 other members of our sample are grouped into a “residual” community (these individuals are likely to be recent immigrants, who are members of communities where most members are outside the United States). Overall, the 81 million users in our primary sample are assigned to 5,140 distinct communities with an average size of 15,910. The average user in our sample has 53.4% of her friends within the same community; at the 10th/50th/90th percentile of our sample, this number is 21%/54%/84%.

timing of these peer effects, showing that an individual acquiring a new phone increases the aggregate propensity that her friends purchase a new phone for at least several months. In Section 3.2, we explore heterogeneities in both influence and susceptibility to influence across demographic characteristics.

Baseline Results. Column 1 of Table 2 presents OLS estimates from regression 6. The results suggest that having one more friend purchase a phone in weeks t or $t - 1$ increases a person’s own propensity to buy a phone in week t by 0.032 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 OLS estimate might also pick up the effects of common shocks or preferences in addition to any peer effects. The rest of Table 2 therefore presents causal effects from IV estimations. 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)
		Broken Phone	Contract Renewal	Broken Phone	Contract Renewal
# of Friends Buying ($t-1$ and t)	0.032 (0.000)	0.046 (0.007)	0.024 (0.014)	0.040 (0.005)	0.022 (0.013)
Controls + Fixed Effects	Y	Y	Y	Y	Y
Mean Dependent Variable	0.93	0.93	0.93	0.93	0.93
Number of Observations	329m	329m	329m	329m	329m
Effective F-Statistic				4,627	878

Note: Table shows estimates of regression 6. 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 corresponding second-stage IV estimates. The dependent variable in all specifications is an indicator for whether user i purchases a new phone in week t . All coefficients reported are multiplied by 100 to ease interpretability. 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. We report Olea and Pflueger (2013) effective F-Statistics. Standard errors are clustered at the level of the community (see the discussion in Section 2.3 and Appendix A.3).

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. This similarity in estimated peer effects across OLS and IV specifications is perhaps surprising, since 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 analyzing the effect of peer purchases on the near-contemporaneous purchasing behavior of individuals.

In terms of magnitudes, a simple back-of-the-envelope calculation suggests that a new phone purchase by an individual in one week leads to an additional 0.08 phone purchases through peer effects in the following week.¹⁴ Put differently, a little less than one in ten phone purchases causes a follow-on purchase in the subsequent week through peer effects.¹⁵

One interesting question is whether the estimated treatment effects are the result of individuals hearing about their friends’ new phone purchase through Facebook or through offline interactions. We think that at most a small part of the overall observed peer effect comes from interactions on Facebook—indeed, in this setting, we view Facebook primarily as a tool to measure phone purchases and social networks, instead of the primary medium for information flow. There are a number of reasons for this. First, only about 2.3% of individuals who post about losing their current phone actually post about purchasing a new phone in the following weeks, and even then, Facebook posts are usually seen by only about a quarter of an individuals’ friends (Bernstein et al., 2013). Second, we highlight below that peer effects from geographically proximate friends are substantially larger, suggesting an important role of in-person interactions in propagating information about new phone purchases. Third, we show below that the effect of a friend’s phone purchase on own purchasing behavior is strong for a number of months following the friends’ purchase. We think it is much more plausible that this effect comes through hearing about the friend’s purchase over time (as well as through second-order peer effects), instead of the delayed effect capturing the purchase of a new phone many months after viewing a social media message—in particular given the evidence that individuals remember only a fraction of social media content even immediately after viewing it (Counts and Fisher, 2011).

The difference in magnitude across the two IV estimates in columns 4 and 5 of Table 2 highlights that the local average treatment effects (LATEs) we capture using each of these instruments may differ from the average treatment effect in the population. Specifically, our first instrument captures the average peer effects of individuals who post publicly about losing their phones (and who then quickly purchase a new one) on those individuals’ friends. Our finding suggests that the peer effects exerted

¹⁴The average peer of people in the sample has 258 friends (which is lower than the average number of friends of people in the sample, which was restricted to only include individuals with at least 100 friends), and a new purchase by these peers increases the probability of each friend purchasing a new phone the following week by about 0.032 percentage points (the average of the two IV estimates). A simple back-of-the-envelope estimate of the overall effect is thus $258 \times 0.00032 \approx 0.08$.

¹⁵We rule out two possible alternative explanations for the patterns in Table 2. First, we explore if they might primarily capture the correlated behavior of family members as a result of contract incentives such as “Buy One, Get One Free” offers that are sometimes available for members of the same family plan. When we repeat our analysis after excluding each user’s family members from their friends (where we identify family members through a combination of self-reports and model-based imputations), we find baseline estimates of very similar magnitude. In addition, while “Buy One, Get One Free” offers might in principle explain correlated purchases that are close in time, they could not explain the long-lasting patterns we show in Section 3.1. We also find that our estimates are not driven by Facebook disproportionately advertising cell phones to people whose friends recently experienced a random phone loss, or whose friends’ contracts were up for renewal. To show this, we repeat our baseline regressions only for users who did not see any cell phone ads on Facebook during our sample period. The peer effects we estimate in this sample are near-identical to those in the full sample. The finding is consistent with our institutional understanding of the scope of ad targeting.

by these individuals may be somewhat larger than the average peer effects in the population, perhaps because individuals who quickly replace a (partially) broken phone care a lot about phones, and are therefore more likely to influence their friends. In 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. In contrast, the IV coefficient estimated using the contract renewal 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, 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 raise the possibility for substantial heterogeneities in peer effects, both along characteristics of the potential influencers and characteristics of the individuals who are potentially influenced.¹⁶ We explore these heterogeneities, which have important implications for firms' marketing strategies and price-setting behaviors, in Section 3.2.


3.1 Peer Effects at Longer Horizons

The specifications reported in Table 2 analyze 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. First, for how long does the purchase of a phone by a peer influence an individual's own purchasing behavior? Second, do these peer effects primarily represent the 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 a user's phone purchasing behavior to include up to 43 weeks following the initial phone purchase by a peer. Specifically, we construct dependent variables of the form $1(\text{BuysPhone}_{i,(t,t+3)})$, $1(\text{BuysPhone}_{i,(t+4,t+7)})$, and so on, to capture whether a user purchases a new phone during a number of four-week periods. In Figure 5, we report the β -coefficients from using these variables as dependent variables in regression 6. Though these regressions are similar to our baseline specification reported in Table 2, the interpretation of the longer-horizon effects is somewhat more complicated. In particular, since individuals and their friends often 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 LATEs associated with a friend purchasing a new phone in weeks t or $t + 1$ on the user purchasing at various horizons, capturing both the direct effect of the initial friend purchase and any higher-order effect of purchases by common friends that were caused by the initial purchase.

¹⁶The 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.

Figure 5: Peer Effects at Alternative Horizons



Note: Figure shows estimates from IV regression 6 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 total 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 IV coefficients in Figure 5. First, 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 probability is even marginally lower in the month prior to the friend's random phone loss, though the effect is barely significant and tiny in magnitude). 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 phone purchase by a friend would behave conditionally similarly in the absence of the random friend purchase.

Second, 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 generally stabilizes. 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 The observed cumulative increase in purchasing probability

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, any such effects are more than offset by positive effects on the number of total purchases through delayed or higher-order peer effects.

is economically meaningful: having an additional friend who purchases a phone in week t increases the chance that a user purchases a phone between weeks t and $t + 15$ by 0.6 percentage points. In our sample, the average chance that a given user purchases a cell phone over this period is 14.8%, so a friend’s purchase increases the user’s own probability of buying a new phone in the next four months by about 4% of the baseline probability.


3.2 Heterogeneities in Treatment Effects

The previous observation 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 on average. 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.4 provides the exact regression specifications. Directionally, the patterns of heterogeneity in the resulting LATEs are generally similar to the patterns of heterogeneity in the corresponding OLS estimates, suggesting that they are not only a feature 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 Strength and Geographic Proximity. We first explore whether the magnitude of the peer effects we observe is affected by the strength of the relationship of the user-friend pair. To measure the closeness of friendship links, we rank a user’s friendships according to a model of tie strength based on characteristics such as mutual friends and interaction frequency, similar to Gilbert and Karahalios (2009). The left panel 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 (we choose this cutoff, since tie strength declines much less strongly across ranks beyond the top 25 friends). It is reassuring that peer effects from closer friends are larger. In fact, there are a number of possible explanations that are consistent with this finding. First, purchases by these friends may be more salient to a user, perhaps because she is more likely to interact with these friends. Second, it is likely that individuals are more willing to trust information that they receive from closer peers. Third, the desire to keep up with closer friends may be higher than the desire to keep up with friends who are less close.

We also explore whether peer effects from geographically proximate friends are larger than those from friends who live further away. The right panel of Figure 6 shows that the estimated peer effect from a friend who lives in the same zip code is more than twice as large as the peer effect from a friend who lives in a different zip code. This evidence is highly consistent with our previous interpretation that much of the observed peer effects are the result of in-person interactions between individuals, which are more likely to occur when two individuals live close to each other than when they live further apart.

Figure 6: Peer Effects Heterogeneity by Relationship Strength and Geographic Proximity




Note: Figure shows IV estimates of equation 6 using the random phone loss instrument. In the left panel, we split each user’s friends into those inside and outside the top 25 using a model of friendship intensity. In the right panel, we split all location-services friends into those living in the same zip code and those living in a different zip code as the user. Error bars show 95% confidence intervals.

Heterogeneity by Friend Characteristics. We next explore heterogeneities 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, and “influencer campaigns” are now an integral part of most consumer marketing strategies (see Ferguson, 2008; Tucker, 2008; Bakshy et al., 2011; Aral and Walker, 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 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 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

Figure 7: Peer Effect Heterogeneity by Friend Characteristics



Note: Figure shows IV estimates of equation 6 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.

per-friend peer effect and relatively many friends. 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 the largest influence on the purchasing behavior of their friends (recall from Table 1 that the median phone age in our sample was 317 days). In turn, individuals who do not regularly replace their phones—and who are therefore likely to not value new technology as much—exert smaller 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.¹⁸

Peer Influence vs. Price Sensitivity. One important 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). The magnitude of this difference depends in part 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.

In our data, we do not have individual-level estimates of price sensitivity. To explore whether the most influential individuals are likely to have relatively high or relatively low price sensitivity,

¹⁸When comparing the 16 IV coefficients presented in the left Panel of Figure 7 to the corresponding coefficients from an OLS specification, we obtain a correlation of 0.81, suggesting that our conclusions regarding the relative strength of peer effects of different individuals may generalize beyond the specific LATE studied here. The main difference is in the heterogeneity by age, where there are fewer differences across age groups in the OLS specification than in the IV specification.

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. 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 2016.¹⁹

We next explore the correlation between peer influence and price sensitivity across the eight groups. We find the correlation with per-friend influence to be 0.89, and the correlation with total influence to be 0.90.²⁰ This result suggests that price cuts disproportionately attract extra demand from individuals who are relatively influential, and that the deviations between individual and aggregate demand curves in this market are thus 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.²¹

The positive correlation between price sensitivity and peer influence may 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, and firms do not want to increase prices to avoid selling to less-influential individuals, 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 can help rationalize, for example, why Apple does not increase the price for its iPhones, despite the large queues outside its stores around device release dates. 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, new restaurants, or the famous Cronuts).


¹⁹On 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. In the week following this price cut, we observe a 4% increase in the number of iPhones registered, with heterogeneity in the size of this jump across demographic groups.

²⁰We obtain similar correlations when we estimate peer effects with OLS regressions (acknowledging the potential biases in these specifications), suggesting the patterns may be generalizable beyond the specific LATE considered here. We 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 0.66 and 0.22 for the per-friend and total peer effects, respectively. Running IV regressions with more endogenous variables is not computationally feasible, preventing us from extending our analysis to consider the correlation between price sensitivity and peer effects at finer demographic splits.

²¹Through this channel, peer effects are a force that lowers markups and improves consumer welfare and allocative efficiency in this market. Appendix A.5 presents a simple model that formally explores this relationship between the correlation of peer influence and price sensitivity, the aggregate demand elasticity, and price markups. It is possible that peer effects may affect optimal price-setting through other channels (see Easley and Kleinberg, 2010; Campbell, 2013; Garcia and Sheling, 2018), and the overall effect of peer effects on prices depends on the relative importance of these various channels.

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. The findings are 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 IV estimates of equation 6 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 area-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 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.

Figure 9: Peer Effect Heterogeneity by Pairwise Characteristics



Note: Figure shows instrumental variables estimates of equation 6 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.

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 observed peer effects are specific to the phone brand purchased by the peer, and explore whether there are positive or negative spillovers to competing brands.

We focus on the two major cell phone lines, Apple’s iPhones and Samsung’s Galaxy phones, which are used by 55% and 27% of users in our data, respectively. We pool the other highly fragmented brands into a residual category, which includes a variety of phones operating largely on the Android system. 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 explore, for example, whether a

friend’s purchase of an iPhone increases a user’s own demand for all phones, including those of iPhone competitor Galaxy, or whether it primarily pulls demand away from Galaxys and towards iPhones.

Identification Challenge and Empirical Approach. To give a concrete example of the challenge with identifying peer effects in phone brand choice, imagine that there are two individuals, Amy and Bob, both of whom have five friends. Among Amy’s friends, four would buy an iPhone if they were to replace their current phones, while only one of Bob’s friends would buy an iPhone. In addition, homophily on characteristics such as tech-savvyness imply that both Amy and Bob are similar to their friends in terms of phone preferences: even in the absence of peer effects, Amy would likely buy an iPhone while Bob would probably buy a different phone. As a result, standard OLS specifications that regress whether people buy a certain phone brand on whether their friends buy that same brand would not necessarily identify peer effects, since correlated preferences (and correlated shocks) would induce similar purchasing behavior even in the absence of any peer effects.

To document the role of peer effects in determining the purchases of specific phone brands, we thus adapt the IV strategy described above. To conceptualize our approach, imagine now that there is a third person, Carl, who is very similar to Amy. Carl also has five friends, out which four would purchase an iPhone if they were to replace their phones, and Carl’s own propensity to purchase an iPhone is also very similar to that of Amy. Now imagine that, in a given week, both Amy and Carl have one of their friends break their phones. By chance, it happens that Amy’s unlucky friend is one that is likely to replace her broken phone with an iPhone, while Carl’s unlucky friend is likely to replace it with a Galaxy. Importantly, this variation in the phone brands bought by Amy’s and Carl’s friends is not driven by differences in the composition of their friends—our thought experiment holds this composition constant by construction. Instead, the brands purchased by the Amy’s and Carl’s friends are determined by which of their friends randomly break their phones in a given week, something that should not be correlated with Amy’s and Carl’s normal purchasing preferences after controlling for the brand preferences of all friends. As a result, any difference in Amy’s and Carl’s probabilities of buying different phone brands in the weeks following their friends’ random phone losses (and subsequent phone replacements) is informative about the causal role of peer effects.

To operationalize this research design, we first construct a measure of each individuals’ probability of purchasing a phone of each brand. Specifically, for each phone brand c , we fit a neural network to predict the propensity that individual i will purchase a phone of brand c in week t , based on observable characteristics of that individual. We predict both the unconditional probability of buying a phone of brand c , $ProbBuyUncond_{i,t}^c$, and the propensity of such a purchase conditional on posting about a random phone loss, $ProbBuyRandomPhoneLoss_{i,t}^c$ (the conditional and unconditional propensities are highly correlated across individuals). We estimate these propensities using information on individuals’ demographics and current phones (see Appendix A.2 for details). For instance, we find that older users prefer iPhones, while all users are more likely to buy a phone of the same brand as their current device. As before, neural networks allow us to uncover non-linear and interactive relationships between the various observable characteristics.

Our research design then proposes to use the sum of these 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 average propensities among all friends. As described above, this is based on the assumption that, conditional on 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. Formally, we calculate, for each individual and each phone brand c , the average conditional probabilities of purchasing phones in each brand category among all her friends as given by equation 8; this will be the central variable to control for the composition of different people’s friends, which could be correlated with those people’s own phone preferences.

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

We also sum up this probability among her friends that randomly lose their phones in a given week as given by equation 9; this will be our instrument for the number of friends purchasing a phone of a particular brand:

$$LossFriendsSumProbBuyRPL_{i,t}^c = \sum_{j \in Fr(i)} \mathbb{1}(RandomPhoneLoss)_{j,t} \cdot ProbBuyRandomPhoneLoss_{j,t}^c. \quad (9)$$


We indeed find homophily in the propensities to buy phones of a certain brand. The left column Figure 10 plots, for each phone brand c , a user’s own $ProbBuyUncond_{i,t}^c$ on the horizontal axis and our instrument for friend purchases of category c , $LossFriendsSumProbBuyRPL_{i,t}^c$, on the vertical axis. We find that individuals who themselves are more likely to purchase a certain brand usually have friends that are also more likely to buy that brand, although the relationships are not always monotonic.

The right column of Figure 10 shows the same relationship as the left column, but conditions on a number of control variables also included in our regressions, the most important of which is $AllFriendsAvgProbBuyRPL_{i,t}^c$. Conditional 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 the brand preferences of person i , at least to the extent that those preferences are captured by observable characteristics such as demographics and current phone brand. This finding makes it more plausible that they are also uncorrelated with brand preferences based on unobservable characteristics of person i , an assumption that is at the heart of our identification strategy.²²

Regression Specification. To study peer effects at the brand level, we 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:

²²As 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 of brand category c , $ProbBuyUncond_{i,t}^c$, on the horizontal axis and the instrument, $LossFriendsSumProbBuyRPL_{i,t}^c$, 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 11, with the exception of $ProbBuyUncond_{i,t}^c$, the variable plotted on the horizontal axis.

$$\begin{aligned} FriendsBuyPhone_{i,(t-1,t)}^{c'} &= \sum_{c \in C} \delta_c^{c'} LossFriendsSumProbBuyRPL_{i,t}^c + \\ &\quad \sum_{c \in C} \phi_c^{c'} AllFriendsAvgProbBuyRPL_{i,t}^c + \omega X_{i,t} + e_{i,t}. \end{aligned} \quad (10)$$

Our three second stages (one for each $c'' \in C$) are of the form:

$$\begin{aligned} \mathbb{1}(BuysPhone)_{i,t}^{c''} &= \sum_{c' \in C} \beta_{c'}^{c''} \widehat{FriendsBuyPhone}_{i,(t-1,t)}^{c'} + \\ &\quad \sum_{c' \in C} \Phi_{c'}^{c''} AllFriendsAvgProbBuyRPL_{i,t}^{c'} + \gamma X_{i,t} + \epsilon_{i,t}. \end{aligned} \quad (11)$$

The indicator variables $\mathbb{1}(BuysPhone)_{i,t}^{c''}$ capture whether person i purchased a phone of brand category 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 central control variable in both the first and second stages of the regression is the average conditional probability of buying a phone of each brand across all of individual i 's friends, $AllFriendsAvgProbBuyRPL_{i,t}^c$. The vector $X_{i,t}$ includes the 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 , given by $ProbBuyUncond_{i,t}^c$, 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.

Since some of the (positive or negative) spillovers across brands would likely materialize only over time, we also 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 3.1, constructing dependent variables of the form $\mathbb{1}(BuysPhone)_{i,(t,t+24)}^{c''}$. We then perform a second set of instrumental variables regressions of the form outlined in Equation 11, replacing the original dependent variables with these multi-period cumulative purchase indicators. As discussed in Section 3.1, the coefficient estimates in these longer-horizon regressions should be interpreted as 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 through purchases of common friends that were induced by this initial purchase.

Estimates of Brand-Level Peer Effects. Table 3 shows results from regression 11. 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. 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.

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	0.027 (0.005)	-0.002 (0.004)	-0.007 (0.004)	0.018 (0.007)	0.340 (0.069)	-0.006 (0.022)	-0.172 (0.043)	0.162 (0.059)
Friends Buy Galaxy	-0.002 (0.009)	0.047 (0.009)	0.019 (0.009)	0.065 (0.016)	-0.335 (0.058)	0.658 (0.047)	0.521 (0.058)	0.844 (0.087)
Friends buy Other	-0.016 (0.007)	-0.012 (0.007)	0.074 (0.009)	0.046 (0.013)	-0.368 (0.051)	0.043 (0.038)	1.229 (0.064)	0.904 (0.079)
Controls + Fixed Effects	Y	Y	Y	Y	Y	Y	Y	Y
Mean Dependent Variable	0.38	0.29	0.25	0.93	11.74	6.44	5.79	23.97
Number of Observations	329m	329m	329m	329m	329m	329m	329m	329m

Note: Table shows estimates of regression 11. 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 users' unconditional probabilities of buying a new phone in each category c , and for the average conditional and unconditional probabilities of purchasing a phone in each category among the users' 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 level of the community (see the discussion in Section 2.3 and Appendix A.3).

We find that friend purchases in each of our three brand categories lead a user to increase their overall probability of purchasing a new phone (see columns 4 and 8). 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. In particular, during our sample period, iPhones were the most well-established brand, suggesting a smaller role for information acquiring through peers; on the other hand, social learning would likely have been most important for the more obscure phones in the fragmented “Other” category.

In addition to these large and positive same-brand peer effects, we also find heterogeneous across-brand demand spillovers. Specifically, we find large positive spillovers from purchases of Samsung Galaxy phones to purchases of phones in the “Other” category; these two brand categories share the Android operating system. This positive demand spillover is also 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. There are fewer spillovers in the other direction, and the small negative spillover from purchases of phones in the “Other” category to purchases of Samsung Galaxy phones is not statistically significant.

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 a user’s demand 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, 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. Second, 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., Shapiro, 2018; Sinkinson and Starc, 2018). Third, positive demand spillovers to other brands using different operating systems could have resulted from perception transfers across competing brands (see Roehm and Tybout, 2006, for related work in the marketing literature). Our finding of substantial negative demand spillovers to competing brands using different operating systems therefore helps researchers understand the implications of peer effects on the competitive dynamics between firms, and distinguish them from the spillover effects of marketing activities.

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 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 generally find positive different-brand demand spillovers for brands sharing an operating system, and 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 about that phone brand, and, to a lesser extent, about other phones using the same operating system. As a result, demand should increase the most for the specific brand purchased by the friend; it should increase somewhat less for competing brands 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.

Peer Effects at the Model Level. In Appendix A.6, we also study peer effects at the device model level, and explore the presence of same-brand, different-model peer effects. 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. For this analysis, we cannot use an instrumental variables research design, as we do in the main body of the paper: while observable characteristics allow us to predict 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. We therefore run OLS specifications that regress an individual’s probability of purchasing a specific phone model on the phone model purchases of her friends. While the absolute magnitudes of the estimates should thus be interpreted with caution, some interesting patterns emerge about the relative size of effects for different phone models. First, same-model peer effects are more than an order of magnitude larger than different-model peer effects. Second, these same-model peer effects do not vary with the cost of the model, but they are decreasing in the time since the model release, providing further evidence for an important social learning channel behind the peer effects. Third, same-brand, different-model peer effects are more than twice as large as different-brand peer effects. The spillovers of peer effects to other models of the same brand are largest for Apple, which co-brands all of its devices under the iPhone brand, and smallest for LG, which does not do so.

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. Our research design cannot precisely identify the channel behind the observed peer effects, but our results are most consistent with an important role of social learning in explaining the observed peer effects. While peer effects expand the overall market for phones, there can be 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 a customer’s value to a firm exceeds the direct effect that this customer has on the firm’s profits.

An interesting question for future work concerns the generalizability of our findings to understanding the decision to adopt products other than cell phones. Indeed, 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 (see, for example, related work by Kuchler, Stroebe, and Wong, 2021). 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 Stroebe, 2015). Specifically, we hope that the increasing availability of social network data, such as the Social Connectedness Index described in Bailey et al. (2018a, 2020a,b,d), will help to improve our understanding of the effects of social interactions on social, political, financial, and economic outcomes.

References

- 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. 2019. "House price beliefs and mortgage leverage choice." *The Review of Economic Studies* 86 (6):2403–2452.
- Bailey, Michael, Patrick Farrell, Theresa Kuchler, and Johannes Stroebel. 2020a. "Social connectedness in urban areas." *Journal of Urban Economics* :103264.
- Bailey, Michael, Abhinav Gupta, Sebastian Hillenbrand, Theresa Kuchler, Robert J Richmond, and Johannes Stroebel. 2020b. "International trade and social connectedness." Tech. rep., National Bureau of Economic Research.
- Bailey, Michael, Drew Johnston, Martin Koenen, Theresa Kuchler, Dominic Russel, and Johannes Stroebel. 2020c. "Social distancing during a pandemic: The role of friends." *NBER Working Paper* (w28234).
- Bailey, Michael, Drew Johnston, Theresa Kuchler, Dominic Russel, Johannes Stroebel et al. 2020d. "The Determinants of Social Connectedness in Europe." In *International Conference on Social Informatics*. Springer, 1–14.
- 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.
- Balachander, Subramanian and Sanjoy Ghose. 2003. "Reciprocal spillover effects: A strategic benefit of brand extensions." *Journal of Marketing* 67 (1):4–13.
- 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.
- Bernstein, Michael S, Eytan Bakshy, Moira Burke, and Brian Karrer. 2013. "Quantifying the invisible audience in social networks." In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 21–30.

- 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.
- Bester, C Alan, Timothy G Conley, and Christian B Hansen. 2011. "Inference with dependent data using cluster covariance estimators." *Journal of Econometrics* 165 (2):137–151.
- Björkegren, Daniel. 2018. "The adoption of network goods: evidence from the spread of mobile phones in Rwanda." *The Review of Economic Studies* 86 (3):1033–1060.
- 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.
- Burszty, 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.
- Campbell, Arthur. 2013. "Word-of-mouth communication and percolation in social networks." *American Economic Review* 103 (6):2466–98.
- 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.
- Conley, Timothy G and Christopher R Udry. 2010. "Learning about a new technology: Pineapple in Ghana." *American Economic Review* 100 (1):35–69.
- Counts, Scott and Kristie Fisher. 2011. "Taking it all in? visual attention in microblog consumption." In *Fifth International AAAI Conference on Weblogs and Social Media*.
- Dahl, Gordon B, Katrine V Løken, and Magne Mogstad. 2014. "Peer effects in program participation." *American Economic Review* 104 (7):2049–74.
- 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.
- 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.
- Easley, David and Jon Kleinberg. 2010. *Networks, crowds, and markets*, vol. 8. Cambridge university press Cambridge.
- Eckles, Dean, René F Kizilcec, and Eytan Bakshy. 2016. "Estimating peer effects in networks with peer encouragement designs." *Proceedings of the National Academy of Sciences* 113 (27):7316–7322.
- 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.
- Facebook. 2016. "Facebook Form 10-Q, Quarterly Report, Q2." .
- 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.
- Garcia, Daniel and Sandro Shelegia. 2018. "Consumer search with observational learning." *The RAND Journal of Economics* 49 (1):224–253.
- 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.
- Gilchrist, Duncan Sheppard and Emily Glassberg Sands. 2016. "Something to talk about: Social spillovers in movie consumption." *Journal of Political Economy* 124 (5):1339–1382.
- 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.
- Goolsbee, Austan and Peter J Klenow. 2002. "Evidence on learning and network externalities in the diffusion of home computers." *The Journal of Law and Economics* 45 (2):317–343.
- 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.
- 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.
- Kuchler, Theresa, Yan Li, Lin Peng, Johannes Stroebel, and Dexin Zhou. 2020. "Social proximity to capital: Implications for investors and firms." Tech. rep., National Bureau of Economic Research.

- Kuchler, Theresa, Dominic Russel, and Johannes Stroebel. 2020. "The geographic spread of COVID-19 correlates with structure of social networks as measured by Facebook." Tech. rep., National Bureau of Economic Research.
- Kuchler, Theresa and Johannes Stroebel. 2020. "Social Finance." Working Paper 27973, National Bureau of Economic Research. URL <http://www.nber.org/papers/w27973>.
- Kuchler, Theresa, Johannes Stroebel, and Arlene Wong. 2021. "Peer Effects and Markups in Product Markets." *Working Paper*.
- 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.
- 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.
- Olea, José Luis Montiel and Carolin Pflueger. 2013. "A robust test for weak instruments." *Journal of Business & Economic Statistics* 31 (3):358–369.
- 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.
- Zacchia, Paolo. 2020. "Knowledge spillovers through networks of scientists." *The Review of Economic Studies* 87 (4):1989–2018.

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 appendix, 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

To construct our primary instrument, we need to identify users who have posted publicly about a random phone loss event. 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 random phone loss events. 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 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 idiosyncratic 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 a natural language processing algorithm.

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%	%cell is broke%	%cell just broke%	
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 $(\overrightarrow{King} - \overrightarrow{Man} + \overrightarrow{Woman})$ is $\overrightarrow{Queen}$.


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 public 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. 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 concerning random phone loss events. 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. 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).¹ 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 $1(\text{RandomPhoneLoss}_{i,t}) = 1$ 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 generates a stronger first stage than was possible with either of the two classifiers used individually.²


¹ROC-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.

²The 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.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, the level of posting 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 *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.

A.2 Purchasing 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),³ 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: Predicting Purchasing Probabilities

	Test Set ROC–AUC	Regression Sample ROC–AUC	Best Layer Sizes
$ProbBuyRandomPhoneLoss_{i, (t, t+1)}$	0.679	0.665	[10]
$ProbBuy2y_{i, (t, t+1)}$	0.640	0.634	[20, 5]
$ProbBuyUncond_{i, t}$	0.705	0.699	[10, 5]
$ProbBuyRandomPhoneLoss^c_{i, (t, t+1)}$	0.720	0.716	[10, 5]
$ProbBuyUncond^c_{i, (t, t+1)}$	0.755	0.752	[20, 5]
$ProbBuyUncond^p_{i, (t, t+1)}$	0.786	0.782	[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 brands or phone models) 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,⁴ current phone model, carrier, age, user browser, Instagram usage flag, state, education level, friend count, 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.

³Training 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.


⁴We 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 the probabilities to purchase a new phone given the random events underlying our two instruments: random phone loss in Panel A, and phone age of 2 years in Panel C. Panels B and D in the right column present Receiver Operating Characteristic (ROC) curves for each of these estimated probabilities. 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 all 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 the probabilities to purchase a new phone of a given brand for individuals with a random phone loss. The right column presents Receiver Operating Characteristic (ROC) curves for each of these instruments. All plots only include users for whom $1(BrokenPhone_{j,t}) = 1$. 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 all observations from the period 2016–18 to 2016–22, the period used to construct our main panel.

In the conditional models, we aim to predict a user’s probability of purchasing a phone in weeks 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 presents information on the performance of the various classifiers for our two baseline instruments. The left column shows binscatter plots of the predicted probability against the realized probability for both a hold-out sample from our training data period (“Pre- and Post-Sample”) and the actual regression data set (“Main Sample”). 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 relatively stable over time. The horizontal axis shows the range of predicted probabilities. For example, we find the predicted 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 the computed conditional probability—rather than just the number of friends with a random phone loss—as an instrument for the number of friends buying a phone.

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 therefore have strong predictive power. We create classifiers at two levels of granularity. First, 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. Second, as described in Appendix A.4, 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: $ProbBuyUncond_{i,t}^c$ for the brand-level granularity, and $ProbBuyUncond_{i,t}^p$ for the model-level granularity. At the brand level, we also train a classifier conditional on a user breaking or losing their phone, $ProbBuyRandomPhoneLoss_{i,t}^c$; these predicted probabilities are used to construct our instruments as described in detail in Section 4, and Figure A.4 explores the performance of these predictors.⁵ Tables A.2 and A.3 show the performance of the brand-level and model-level propensity predictions.

⁵We do not estimate IV regressions at the model level, since we find that the predictions within a brand tend to be highly collinear, making first stage estimation complicated (see Appendix A.4 for further discussion).

Table A.3: Phone-Specific ROC–AUC

	Unconditional		Conditional on Losing Phone	
	Test Set ROC–AUC	Regression Sample ROC–AUC	Test Set ROC–AUC	Regression Sample 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 (see also Table A.2). 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 A.6.

A.3 Statistical Inference – Robustness Exercises

As discussed in Section 2.3, to construct our baseline standard errors, we follow the suggestions by Eckles, Kizilcec, and Bakshy (2016) and Zacchia (2020) and partition the Facebook social graph into a number of communities with limited cross-community dependence, allowing us to cluster standard errors at the community level. Starting from a sparse matrix representing the Facebook social graph, a distributed variant of the Kernighan-Lin algorithm is used to divide the global Facebook social graph into about 21,000 distinct communities. Individuals in our sample are assigned to their communities created by this algorithm. The 0.2% of our sample assigned to communities with fewer than 100 other members of our sample are grouped into a “residual” community (these individuals are likely to be recent immigrants, who are members of communities where most members are outside the United States). Overall, the 81 million users in our primary sample are assigned to 5,140 distinct communities with an average size of 15,910 users. The average user in our sample has 53.4% of her friends within the same community; at the 10th/50th/90th percentiles of our sample, these numbers are 21%/54%/84%.

The first row of Figure A.5 compares the standard errors using this clustering approach to baseline heteroskedasticity-robust standard errors for the estimates corresponding to column 4 (left panel) and column 5 (right panel) of Table 2. We find that the community-clustered standard errors are very similar in magnitude to the heteroskedasticity-robust standard errors, suggesting that, in our setting, across-individual spillovers seem to not confound our inference, at least after conditioning on the large set of controls and fixed effects. The other two rows compare standard errors clustered at the county level (middle row) and individual level (bottom row) to heteroskedasticity-robust standard errors. As before, these standard errors are at most 7.2% larger (and, in some cases, even marginally smaller) than heteroskedasticity-robust standard errors.

Figure A.5: Comparison of Standard Errors



Note: Figure compares standard errors using various clustering approaches to heteroskedasticity-robust standard errors for the instrumental variables estimates corresponding to column 4 (left panel) and column 5 (right panel) of Table 2. The top row uses standard errors clustered at the community level (see Appendix A.3 for discussion on our approach to detecting communities), the middle row uses standard errors clustered at the county level, and the bottom row uses standard errors clustered at the user level.

A.4 Heterogeneity Models

In Section 3.2, 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 strong and weak friendships, and another to geographically proximate or non-proximate friends. 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 $g \in G$:

$$\begin{aligned} Instrument_{i,t}^{Lose,g} &= \sum_{j \in Fr(i)} ProbBuyRandomPhoneLoss_{j,t} \cdot \mathbb{1}(LostPhone_{j,t}) \cdot \mathbb{1}(Condition_{j,t}^g) \\ FriendsBuyPhone_{i,(t,t+1)}^g &= \sum_{j \in Fr(i)} \mathbb{1}(BuysPhone)_{j,(t,t+1)} \cdot \mathbb{1}(Condition_{j,t}^g) \end{aligned}$$

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:

$$\begin{aligned} Friends_{i,t}^g &= \sum_{j \in Fr(i)} \mathbb{1}(Condition_{j,t}^g) \\ AllFriendsAvgProbBuyRPL_{i,t}^g &= \frac{1}{Friends_{i,t}^g} \sum_{j \in Fr(i)} ProbBuyRandomPhoneLoss_{j,t}^g \cdot \mathbb{1}(Condition_{j,t}^g) \end{aligned}$$

Using these new variables, we estimate one first stage per condition as well as a single second stage:

$$\begin{aligned} FriendsBuyPhone_{i,(t-1,t)}^g &= \sum_{g \in G} \delta_g Instrument_{i,t-1}^{Lose,g} + \sum_{g \in G} \psi_g Friends_{i,t-1}^g + \\ &\quad \sum_{g \in G} \theta_g AllFriendsAvgProbBuyRPL_{i,t-1}^g + \omega X_{i,t} + e_{i,t} \\ \mathbb{1}(BuysPhone_{i,t}) &= \sum_{g \in G} \beta_g \widehat{FriendsBuyPhone}_{i,(t-1,t)}^g + \sum_{g \in G} \Psi_g Friends_{i,t-1}^g + \\ &\quad \sum_{g \in G} \Theta_g AllFriendsAvgProbBuyRPL_{i,t-1}^g + \gamma X_{i,t} + \epsilon_{i,t} \end{aligned}$$

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

Heterogeneity by User Characteristics. In Section 3.2, 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:

$$Instrument_{i,t}^{Lose,g} = \mathbb{1}(Condition_{i,t}^g) \cdot \sum_{j \in Fr(i)} ProbBuyRandomPhoneLoss_{j,t} \cdot \mathbb{1}(LostPhone_{j,t})$$

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

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. As before, we have one first stage for each group g , as well as one second stage regression:

$$\begin{aligned} FriendsBuyPhone_{i,(t-1,t)}^g &= \sum_{g \in G} (\delta_g Instrument_{i,t-1}^{Lose,g}) + \omega X_{i,t} + e_{i,t} \\ \mathbb{1}(BuysPhone_{i,t}) &= \sum_{g \in G} (\beta_g \widehat{FriendsBuyPhone}_{i,(t-1,t)}^g) + \gamma X_{i,t} + \epsilon_{i,t} \end{aligned}$$

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

Heterogeneity by Pairwise Characteristics In Section 3.2, 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 friends. We construct new instruments and endogenous variables as follows:

$$\begin{aligned} Instrument_{i,t}^{Lose,g_1,g_2} &= \mathbb{1}(Condition_{i,t+1}^{g_1}) \cdot \\ &\quad \sum_{j \in Fr(i)} [ProbBuyRandomPhoneLoss_{j,t} \cdot \mathbb{1}(LostPhone_{j,t}) \cdot \mathbb{1}(Condition_{j,t}^{g_2})] \\ FriendsBuyPhone_{i,(t,t+1)}^{g_1,g_2} &= \mathbb{1}(Condition_{i,t+1}^{g_1}) \cdot \sum_{j \in Fr(i)} [\mathbb{1}(BuysPhone)_{j,(t,t+1)} \cdot \mathbb{1}(Condition_{j,t}^{g_2})] \end{aligned}$$

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

$$\begin{aligned} FriendsBuyPhone_{i,(t-1,t)}^{g_1,g_2} &= \sum_{g_1 \in G} \sum_{g_2 \in G} (\delta_{g_1,g_2} Instrument_{i,t-1}^{Lose,g_1,g_2}) + \sum_{g \in G} \psi_g Friends_{i,t-1}^g + \\ &\quad \sum_{g \in G} \theta_g AllFriendsAvgProbBuyRPL_{i,t-1}^g + \omega X_{i,t} + e_{i,t} \\ \mathbb{1}(BuysPhone_{i,t}) &= \sum_{g_2 \in G} \sum_{g_1 \in G} (\beta_{g_1,g_2} \widehat{FriendsBuyPhone}_{i,(t-1,t)}^{g_1,g_2}) + \sum_{g \in G} \Psi_g Friends_{i,t-1}^g + \\ &\quad \sum_{g \in G} \Theta_g AllFriendsAvgProbBuyRPL_{i,t-1}^g + \gamma X_{i,t} + \epsilon_{i,t} \end{aligned}$$

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.

A.5 Theoretical Model: Peer Effects, Demand Elasticities, and Prices

In this appendix, we describe a simple model of price setting under monopolistic competition that allows us to illustrate an important channel through which peer effects in consumption can affect demand elasticities and markups.⁶

A.5.1 Consumer Preferences

There are N consumers, and the consumption of each consumer is infinitesimally small relative to total demand. An individual consumer i 's utility function is given by:

$$U_i = \left(\int_0^n R_i(Q_j) q_{ij}^{\rho_i} dj \right)^{\frac{1}{\rho_i}}, \quad (\text{A.1})$$

where q_{ij} is individual i 's consumption of variety j , $Q_j = \{q_{1j}, q_{2j}, \dots, q_{Nj}\}$ is a vector of quantities consumed by all other individuals, and n is the mass of varieties available to consumers. The parameter ρ_i is a measure of substitutability across product varieties, and is allowed to vary across individuals.

The function $R_i(\cdot)$ provides a reduced-form way of capturing the dependence of a consumer's consumption utility on the consumption of others through peer effects. $R_i(\cdot)$ can be individual-specific, so that the quantity consumed by others enter differentially into different consumers' utility; in other words, different individuals are allowed to be differentially affected by the consumption of their peers.

One channel through which peers' consumption can influence an individual's own utility from a particular good is a desire of people to consume similar types of products as others. This could come about because of the proverbial "keeping up with the Joneses," or because of network externalities that make a particular good more useful if others are also using it (e.g., being on the iOS operating system is more useful to me if my friends are also on iOS, since we can then communicate using Facetime). In this case, the function R would be positive and increasing in q_{-ij} , ensuring that individual i 's utility for consuming a product is higher when the product is also widely consumed by others.

An alternative channel through which an individual's own consumption of goods may be affected by the consumption of others is through social learning. For instance, suppose that a product only enters the consumer's choice set if they learn about the existence of the item from friends that already bought the item. In this case, we could specify R to reflect the probability that the individual learns about the item from their friends, where R is an increasing function of the purchases made by friends. At the extreme, if the individual has one friend and she only learns about the item if the friend purchased the item, then $R(q_{-ij}) = 1$ when $q_{-ij} > 0$, and $R(q_{-ij}) = 0$ otherwise.

Parametric example. For the purposes of illustrating the implications of peer effects for consumption and markups, consider the following parametric form:

$$R_i(Q_j) = \left(\int_{F(i)} \psi_f \times q_{fj} df \right)^{\eta_i}, \quad (\text{A.2})$$

⁶As discussed in the paper, it is possible that peer effects can also influence prices and markups through other channels. In that case, the overall effect will depend on the relative strength of the forces discussed here, and any additional mechanisms.

so that the utility for individual i is given by

$$U_i = \left(\int_0^n q_{ij}^{\rho_i} \left(\int_{F(i)} \psi_f \times q_{fj} df \right)^{\eta_i} dj \right)^{\frac{1}{\rho_i}}, \quad (\text{A.3})$$

where $F(i)$ denotes the set of peers that influence individuals i , and q_{fj} denotes the quantity of variety j consumed by individual i . The two scalar variables η_i and ψ_f affect how sensitive individual i 's consumption is to the consumption of her peers. η_i reflects how susceptible individual i is to the peer effects exerted by others. When $\eta_i = 0$, then there are no peer effects from others and we have the usual CES utility function. When $\eta_i > 0$, then an individual's utility is dependent on the consumption of their peers. The scalar ψ_f reflects how influential individual f 's consumption for the utility and consumption of f 's friends. Larger values for ψ_f mean the consumption of individual f has a greater impact on the utility of her friends.

Together, η_i and ψ_f affect the overall magnitude of the peer effects in driving individual i 's consumption. We can see this in the elasticity of expression A.2 with respect to quantity consumed by friend f :

$$\frac{\partial \ln R_i(Q_j)}{\partial \ln q_{fj}} = \eta_i \tilde{\psi}_{fj}, \quad (\text{A.4})$$

where

$$\tilde{\psi}_{fj} = \frac{\psi_f q_{fj}}{\int_{F(i)} \psi_g q_{gj} dg}. \quad (\text{A.5})$$

The higher the value of η_i and ψ_f , the more sensitive is individual i 's consumption of variety j to changes in peer f 's consumption of variety j .

We next derive the analytical solutions for prices and markups using the parametric specification in equation A.2. However, the implications for elasticities of consumption and markups hold for various forms of $R_i(Q_j)$ that have the property where $R_i(Q_j)$ is weakly increasing in Q_j .

A.5.2 Consumer Demand

Consumer i 's constrained maximization problem is given by

$$\max_{q_{ij}} U_i^{\rho_i} - \lambda_i \left(\int_0^n p_j q_{ij} dj - I_i \right),$$

where λ_i is the multiplier on the budget constraint, and I_i denotes total income. From the consumer's maximization problem, the Frisch demand function is given by

$$q_{ij} = \left(\frac{\lambda_i p_j}{\rho_i \left(\int_{F(i)} \psi_f \times q_{fj} df \right)^{\eta_i \rho_i}} \right)^{\frac{1}{\rho_i - 1}}. \quad (\text{A.6})$$

The relative demand for varieties j and k is given by

$$\frac{q_{ij}}{q_{ik}} = \left(\frac{p_j \left(\int_{F(i)} \psi_f \times q_{fk} df \right)^{\eta_i \rho_i}}{p_k \left(\int_{F(i)} \psi_f \times q_{fj} df \right)^{\eta_i \rho_i}} \right)^{\frac{1}{\rho_i - 1}}. \quad (\text{A.7})$$

Let $\sigma_i \equiv \frac{1}{1-\rho_i}$, which is the elasticity of substitution when $\eta_i = 0$. Hence,

$$q_{ij} = q_{ik} \left(\frac{p_j}{p_k} \right)^{-\sigma_i} \left(\frac{\int_{F(i)} \psi_f \times q_{fk} df}{\int_{F(i)} \psi_f \times q_{fj} df} \right)^{\eta_i (1-\sigma_i)}. \quad (\text{A.8})$$

Multiplying both sides by p_j and integrating with respect to all varieties yields

$$\int_0^n p_j q_{ij} dj = \int_0^n q_{ik} p_j^{1-\sigma_i} p_k^{\sigma_i} \left(\frac{\int_{F(i)} \psi_f \times q_{fk} df}{\int_{F(i)} \psi_f \times q_{fj} df} \right)^{\eta_i (1-\sigma_i)} dj.$$

The left-hand side is total expenditure, which we assume equals total income I_i . Hence, we can write the individual consumer's demand for variety k as

$$q_{ik} = \frac{p_k^{-\sigma_i} \left(\int_{F(i)} \psi_f \times q_{fk} df \right)^{\eta_i (\sigma_i - 1)}}{\int_0^n p_j^{1-\sigma_i} \left(\int_{F(i)} \psi_f \times q_{fj} df \right)^{\eta_i (\sigma_i - 1)} dj} I_i. \quad (\text{A.9})$$

Elasticity of Consumer Demand

The individual consumer's demand for variety k (and the individual's price elasticity of demand) now depends on the consumption of others. From equation A.8, we can compute the elasticity of an individual consumer i 's demand for variety k relative to the demand of others (e.g., friend f). This is given by

$$\frac{\partial \ln q_{ik}}{\partial \ln q_{fk}} = (\sigma_i - 1)(1 - s_{ik})\eta_i \tilde{\psi}_{fk} \quad (\text{A.10})$$

where q_{ik} is the quantity consumed by individual i , q_{fk} is the quantity consumed by individual f , $Q_j = \{q_{1j}, q_{2j}, \dots, q_{Nj}\}$ is a vector of quantities consumed across individuals, and s_{ik} is the expenditure share of variety k for consumer i .

As discussed above, the scalar η_i reflects the susceptibility of individual i to peer effects, and $\tilde{\psi}_{fj}$ is the influence of friend f defined in equation A.5. The larger is η_i , the more sensitive is the individual's demand for variety j to the consumption of others; similarly, increases in consumption by more influential friends (high- $\tilde{\psi}_{fj}$ friends) have a larger effect on own demand for variety j . When $\eta_i = 0$, the individual's demand does not depend on consumption of others.

Price elasticity of demand

The price elasticity of consumer i 's demand for variety k is given by:⁷

$$\frac{-\partial \ln q_{ik}}{\partial \ln p_k} = \sigma_i + (\sigma_i - 1)\eta_i \int_{F(i)} \tilde{\psi}_{fk} \times \frac{-\partial \ln q_{fk}}{\partial \ln p_k} df. \quad (\text{A.11})$$

Equation A.11 highlights that, for the standard case with $\sigma > 1$, peer effects lead individuals to have larger price elasticities than in the benchmark model (which only includes the first term). In addition, the equation shows that individuals can have different price elasticities of demand for at least three reasons.

1. They may have different elasticities of substitution between varieties, i.e. heterogeneity in σ_i . The larger σ_i , the more responsive is consumer i 's demand to price changes. This effect, which corresponds to the first term in equation A.11, exists even in the absence of peer effects.
2. Individuals can have different price elasticities because their utility is differently affected by the consumption of others. In other words, individuals may be differentially susceptible to peer effects, generating heterogeneity in η_i . The larger is η_i , the more price elastic is the demand of individual i , all else equal, because the increase in consumption of their peers in response to a given price drop has a larger effect on their own utility (and thus demand).
3. Individuals may also have heterogeneous price elasticities of demand because their set of friends is different in terms of both their influence and their price elasticities. Specifically, peer effects can lead to higher price elasticity of demand when (i) a person's friends are more influential (the first term within the integral, $\tilde{\psi}_{fk}$), and (ii) for a given $\tilde{\psi}_{fk}$, the consumption of peers is more price sensitive (the second term within the integral, $\frac{-\partial \ln q_{fk}}{\partial \ln p_k}$). Putting the two forces together, we can see that there is a larger price elasticity of demand when each of these two terms is higher, as well as when there is a higher covariance between the peer influence of a friend and the price sensitivity of the friend. Our empirical micro estimates are informative about the magnitudes of the covariance between the peer influence of a friend and the price elasticity of that friend.

A.5.3 Firm's problem and Price Setting

Production of a good involves a fixed cost F in addition to a constant marginal cost c , so the average cost is decreasing in quantity. To keep things simple, there are no economies of scope to producing multiple varieties. Therefore, there is a continuum of firms, where each firm produces one variety and there is one firm per variety. A firm that produces variety k has profits

$$\pi_k = p_k x_k - c x_k - F,$$

⁷This derivation is based on the assumption that there is a continuum of firms and varieties produced, and the function R is bounded. Therefore, $\frac{\partial \int_0^n p_j^{1-\sigma_i} R_i(Q_j)^{(\sigma_i-1)} dj}{\partial p_k} \approx 0$.

where x_k denotes the quantity of variety k produced. The firm sets prices p_k to maximize profits. Solving the firm's problem yields:

$$p_k = c + \frac{-x_k}{\partial x_k / \partial p_k}. \quad (\text{A.12})$$

In equilibrium, the free entry condition implies zero profits. Hence, $x_k = F/(p_k - c)$. The market clearing condition is given by $x_k = \int_i q_{ik} di$, where q_{ik} is the quantity of variety k consumed by individual i . From the market clearing condition, we can derive the following partial derivative expression:

$$\frac{-\partial x_k}{\partial p_k} \frac{1}{x_k} = \int_i \frac{-\partial q_{ik}}{\partial p_k} di \times \frac{1}{\int_f q_{fk} df} = \frac{1}{p_k} \int_i \frac{-\partial \ln q_{ik}}{\partial \ln p_k} w_{ik} di, \quad (\text{A.13})$$

where $w_{ik} = \frac{q_{ik}}{\int_f q_{fk} df}$, and the price elasticity of demand $\frac{\partial \ln q_{ik}}{\partial \ln p_k}$ was previously derived in equation A.10.

Equation A.13 can be interpreted as the weighted-average aggregate price elasticity of demand, where the individual price elasticities of demand are weighted by their quantity of consumption relative to total aggregate consumption. Substituting A.13 into the firm's first-order conditions gives the following markup expression:

$$\frac{p_k}{c} = \frac{1}{1 - \theta_k}, \quad (\text{A.14})$$

where

$$1/\theta_k \equiv \int_i \left(\sigma_i + (\sigma_i - 1)\eta_i \int_{F(i)} \tilde{\psi}_{fk} \times \frac{-\partial \ln q_{fk}}{\partial \ln p_k} df \right) w_{ik} di \quad (\text{A.15})$$

Notice that when there are no consumption peer effects, i.e. $\eta_i = 0$, then the markup equals $\frac{\sigma_k}{\sigma_k - 1}$, where $\sigma_k \equiv \int_i \sigma_i w_{ik} di$. This corresponds to the markup expression in the usual monopolistic competition Dixit-Stiglitz model, allowing for heterogeneity in σ across individuals.

In the presence of consumption peer effects, where $\eta_i > 0$, markups are strictly less than $\frac{\sigma_k}{\sigma_k - 1}$. Peer effects are stronger (and markups are smaller) when individuals' utilities of a variety is more sensitive to the consumption of others (large η_i). Markups are also smaller when price changes lead to larger changes in the amount of peer influence exerted, either because each friend is more influential, or because each friend is more price sensitive (and therefore more likely to buy and exert influence), or because the most price sensitive friends are the most influential ones. Equation A.15 also highlights another covariance that determines the overall influence of peer effects on markups: the correlation between an individuals' own susceptibility to peer influence (η_i) and the overall peer effect exerted on this individual, which depends on which people are friends with that individual (F_i): markups are particularly low if the most susceptible individuals are friends with the most influential individuals.⁸

Equation A.14 has also has implications for the dispersion of markups across firms and over time. In particular, it shows that markups can vary across firms because of differences in the strength of peer effects for different products produced by firms. Indeed, our estimates imply that the strength of peer effects do vary significantly across different mobile phone brands, and there are likely even stronger heterogeneities across different product categories that could contribute to markup dispersion. In related work, we explore this further (see Kuchler, Stroebe, and Wong, 2021)

⁸In our setting, we did not find a large amount of heterogeneity in susceptibility to peer influence across demographic groups, but these parameters will differ in other settings.

A.6 Phone Model-Level Analysis

In this appendix, 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 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, as we do in the main body of the paper. 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.⁹ Instead of an instrumental variables specification, we therefore run the following OLS regressions for each model $p' \in P$:


$$\begin{aligned} \mathbb{1}(BuysPhone)_{i,t}^{p'} = & \sum_{p \in P} \beta_p^{p'} FriendsBuyPhone_{i,t-1}^p + \\ & \sum_{p \in P} \pi_p^{p'} ProbBuyUncond_{i,t}^p + \\ & \sum_{p \in P} \zeta_p^{p'} AllFriendsAvgProbBuyUncond_{i,t}^p + \gamma X_{i,t} + \epsilon_{i,t} \end{aligned} \quad (\text{A.16})$$

The central controls in these specifications are the unconditional probabilities of purchasing each phone $p \in P$ of person i , as well as the averages of these probabilities across all of person i ’s friends. The vector $X_{i,t}$ includes additional controls analogous to those in Equation 11. The coefficients of interest are the parameters $\beta_p^{p'}$, of which there are at total of $|P| \times |P| = 529$. The similarity of the OLS and IV estimates in Section 3 suggests that biased coefficients due to correlated shocks may not be a first-order concern for the analysis, especially after conditioning on time fixed effects interacted with our large set of control variables. In addition, even if these OLS estimates were biased on average, patterns across these coefficients can still be informative about the nature of the underlying peer effects—especially since many common shocks should affect these estimates in similar ways.

The left panel of Figure A.6 shows a histogram of the estimated $\beta_p^{p'}$ coefficients. About 64% of the estimated coefficients are positive. The right panel shows the distributions of $\beta_p^{p'}$ coefficients for three groups. The 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.

⁹In Appendix A.2, we describe how we estimate the unconditional probabilities that individuals purchase a phone of each type p in a given week, $ProbBuyUncond_{i,t}^p$. 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.

Figure A.6: Specific Model Peer Effects (OLS)




Note: Figure shows the distributions of $\beta_p^{p'}$ coefficients from regression A.16. 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. We 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 for phones 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. The box plots show the 5th, 25th, 50th, 75th and 95th percentiles of the distribution.

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 A.7, 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, though, as we have discussed, we cannot rule out that it also contributes to the overall peer effects.

Figure A.6 also shows that same-brand, different-model peer effects are almost three times as large as different-brand peer effects, with an average $\beta_p^{p'}$ coefficient of 0.0036 vs. 0.0013. To further explore the same-brand, different-model peer effects, we split them up according to the three major brands in our data. The left panel of Figure A.8 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


Figure A.7: Same-Model Peer Effects



Note: Figure shows scatterplots of the 20 same-model $\beta_p^{p'}$ coefficients from regression A.16 against the price of the model (left panel) and the time since the model's release (right panel).

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, for academic 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 by Balachander and Ghose (2003).


Figure A.8: Same-Brand, Different-Model Peer Effects



Note: Figure shows the same-brand, different-model $\beta_p^{p'}$ coefficients from regression A.16. 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 A.8 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

Figure A.9: Different-Brand Peer Effects



Note: Figure shows the different-brand $\beta_p^{p'}$ coefficients from regression A.16. 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.

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 across-product spillovers decrease in magnitude as products become more dissimilar (e.g., Janakiraman, Sismeiro, and Dutta, 2009).


In the last part of the analysis, we split out different-brand peer effects. The left panel of Figure A.9 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 take-aways result from our model-level analysis. First, same-model peer effects are more than an order of magnitude larger than different-model peer effects. Second, 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. Third, 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 are largest for models of similar value.

A.7 Robustness Checks – Specifications

In our baseline regressions, we include a number of controls and fixed effects, which we describe in detail in Section 2.3. In this appendix, we show that our results are robust to changes in the set of controls and fixed effects we incorporate. In Figure A.10, we present the estimated β -coefficients for 64 different variants of our baseline regression 6. Each variant is depicted as a dot, with the 95% confidence intervals shaded in gray. The red diamond corresponds to our baseline specification. In the panel below the graph, we indicate which fixed effects and controls are included in each regression. We consider three modifications to the set of fixed effects. In the regressions marked *No_Week*, we omit the week interaction term from each of the three sets of fixed effects. Correspondingly, in the regressions marked *No_Edu*, we remove the education fixed effect, while in the regressions marked *No_Age*, the age fixed effect is omitted. For each of the four variants of fixed effects (our baseline plus three adjustments), we run a number of permutations, excluding some of our baseline set of controls. In the regressions marked *User_ML*, we control for if a user posts about breaking their own phone. In the regressions marked *Friend_Privacy*, we control for the number of a user's friends whose posts are public by default. In the regressions marked *Other_Posts*, we control for the number of a user's friends who post about anything other than a broken phone. In the regressions marked *Historic_Prob*, we control for the number of a user's friends who broke their phone in the 12 months prior to our sample, or for the conditional probability of buying a new phone among these users. In our baseline specification, we control for all of these variables, but our results are largely similar if we only control for a subset of them or if we adjust the fixed effects.

Figure A.10: Robustness of Alternative Specifications of Controls and Fixed Effects



Note: Figure compares estimates of the β -coefficients for 64 different variants of our baseline regression 6. Our baseline specification is marked with a diamond; the panel below the graph indicates the specification corresponding to each