The Economic Effects of Social Networks: Evidence from the Housing Market

Michael Bailey

Facebook

Ruiqing Cao

Harvard University

Theresa Kuchler

New York University

Johannes Stroebel

New York University, National Bureau of Economic Research, and Centre for Economic Policy Research

We show how data from online social networking services can help researchers better understand the effects of social interactions on economic decision making. We combine anonymized data from Facebook, the largest online social network, with housing transaction data and explore both the structure and the effects of social networks. Individuals whose geographically distant friends experienced larger recent house price increases are more likely to transition from renting to owning. They also buy larger houses and pay more for a given house. Survey data show that these relationships are driven by the effects of social interactions on individuals' housing market expectations.

This paper was previously circulated as "Social Networks and Housing Markets." For helpful comments, we are grateful to Alex Chinco, Tony Cookson, Eduardo Davila, Anthony DeFusco, Giovanni Favara, Andreas Fuster, Marty Eichenbaum, Xavier Gabaix, Pedro Gete, Stefano Giglio, Adam Guren, Ben Golub, Harrison Hong, Erik Hurst, Ben Keys, Andres

Electronically published October 30, 2018
[*Journal of Political Economy*, 2018, vol. 126, no. 6]

Researchers are increasingly interested in studying the effects of interactions through social networks on economic decision making. However, analyzing the economic effects of social interactions has proved challenging, in large part because of the absence of high-quality data on social networks that can be linked to outcome variables of interest. In this paper, we show how data from online social networking services such as Facebook and LinkedIn can help overcome this measurement challenge, with the potential to dramatically expand our understanding of the role of social networks across a large number of settings. We illustrate this point by using anonymized social network data from Facebook to analyze the role of social interactions in the housing market. We show that the recent house price experiences within an individual's social network affect her perceptions of the attractiveness of property investments and through this channel have large effects on her housing market investments.

We observe an anonymized snapshot of the "social graph" of friendship links on Facebook. Facebook is the world's largest online social network, with over 234 million active users in the United States and Canada and more than 1.9 billion users globally. We argue that social networks as measured by Facebook provide a realistic representation of real-world US friendship networks. As we discuss below, this is the result of Facebook's enormous scale, the relative representativeness of its user body, and the fact that people primarily use Facebook to interact with their real-world friends and acquaintances.

We begin by documenting salient features of the observed US friendship networks, with a focus on elements of social network structure that have been linked to social and economic phenomena such as the diffusion of information and the construction of social norms. There is significant across-individual variation in both network size and local network clustering (the probability of two friends of an individual being friends with each other). Network size declines in age, while local clustering is U-shaped in age, with the oldest individuals having the smallest and most

Liberman, Michael Luca, Brigitte Madrian, Ulrike Malmendier, Holger Mueller, Stijn van Nieuwerburgh, Christopher Palmer, Cecilia Parlatore, Alp Simsek, Andrei Shleifer, Joe Vavra, Arlene Wong, Wei Xiong, Basit Zafar, and Eric Zwick, four anonymous referees, as well as seminar and conference participants at Baruch, the Consumer Financial Protection Bureau, Conference on Housing, Urban, and Macroeconomics, Chicago, Dartmouth Tuck, Financial Conduct Authority, the Federal Reserve Banks of San Francisco, Philadelphia, and New York, Frankfurt, German Economists Abroad Christmas Conference, Harvard, Harvard Business School, London School of Economics, Mannheim, Minnesota, National Bureau of Economic Research (Corporate Finance, Behavioral Macro, Behavioral Finance, Impulse and Propagation, Digitization, Household Finance, Real Estate), Northwestern, New York University, Penn State, Society for Economic Dynamics, Stanford Institute for Theoretical Economics, Stanford Graduate School of Business, University of California Berkeley and Davis, and University of British Columbia. We thank the Center for Global Economy and Business at NYU Stern for generous research support. This research was facilitated through a research consulting agreement between the academic authors and Facebook.

clustered networks. More educated individuals have larger and less clustered networks. Networks of urban and rural individuals are relatively similar in size, though rural networks are somewhat more clustered. Despite these systematic patterns, most of the across-individual variation in network characteristics occurs within rather than across these demographic groups. We also document significant homophily, with individuals more likely to be friends with others that are similar on demographic and network characteristics. When we explore the geographic dimension of the US friendship network, we find that while the average person has friends in 71 different counties, more than 34 percent of her friends live in the same county, and 63 percent live in the same state. Similarly, for the average person, 53 percent of friends live within 50 miles, and 67 percent live within 200 miles. The geographic concentration of friendship networks is declining in both age and education and varies significantly across counties in the United States.

In the second part of the paper, we exploit the across-individual variation in the structure of social networks to analyze the effects of social interactions on individuals' housing investment decisions. To measure housing investment decisions, we combine the social network information from Facebook with anonymized public-record data on individuals' housing transactions for Los Angeles County. Our final sample contains anonymized data on 1.4 million individuals and 525,000 housing transactions. We use these combined data to analyze the effects of the house price experiences within an individual's social network on three aspects of her housing market investment behavior: the extensive margin decision (i.e., whether to rent or own), the intensive margin decision (i.e., the square footage of properties bought), and the willingness to pay for a particular house.

Our analysis starts by documenting that, at any point in time, different people in the same local housing market have friends who have experienced vastly different recent house price movements. For example, the average 2008–10 house price changes experienced by the friends of the individuals in our Los Angeles sample ranged from −10.1 percent at the 5th percentile of the distribution to −5.2 percent at the 95th percentile of the distribution. This variation is driven by heterogeneity across our sample in the location of individuals' friends, combined with variation in regional house price changes.

In order to isolate a causal relationship between friends' house price experiences and own housing market investments, we need to rule out potential noncausal explanations of any observed correlation. A first concern is that our interpretation could be confounded if individuals form expectations by extrapolating from their own house price experiences, which would be correlated with the house price experiences of their local friends. To exploit only variation in friends' house price experiences that is orthogonal to a person's own experiences, we instrument for the house price experiences of all friends with the experiences of only her friends in geographically distant housing markets.

Using this instrumental variables strategy, we show that the house price experiences within an individual's social network have a quantitatively large effect on all three aspects of her housing investment decision. First, a 5 percentage point higher average house price change between 2008 and 2010 in the counties where an individual has friends leads to a 3.0 percentage point increase in the probability of that individual transitioning from being a renter in 2010 to being a homeowner in 2012, relative to a baseline transition probability of 18 percent. This is more than half the effect size of adding a family member. We also find that homeowners are more likely to transition to renting when their friends experience below-average house price changes. Second, conditional on an individual buying a house, a 5 percentage point increase in friends' house price experiences over the 24 months prior to the purchase is associated with the individual buying a 1.6 percent larger property. Third, conditional on observable property characteristics, a 5 percentage point increase in the house price experiences in an individual's social network is associated with that individual paying 2.3 percent more for the same property. This estimate is robust to adding property fixed effects to account for unobservable property characteristics. When we also control for the house price movements in the seller's network, we find that sellers whose friends had more positive house price experiences also demand higher sales prices.

We argue that these relationships between the house price experiences in an individual's social network and her housing market behavior capture a causal mechanism. In addition to using our instrumental variables strategy to abstract from a possible extrapolation of own house price experiences, we rule out a number of alternative noncausal explanations of our estimates.

In particular, we address possible challenges coming from the nonrandom exposure of individuals to different geographically distant housing markets. We first show that a correlation between where an individual has friends and her own characteristics does not, by itself, confound our findings. The reason is that the house price experiences within an individual's social network are affected by the interaction of the geographic distribution of her friends and how house prices in those areas move in a given year. While people with friends in Boston are different from people with friends in Miami, relative house price movements in Boston and Miami change over time. Comparing the housing investment behavior of individuals with friends in Boston across different years thus removes any time-invariant confounding effect of the geographic distribution of an individual's friends. In fact, in some specifications we observe multiple transactions of the same individual across different years. We find that this same individual is willing to pay more for a given house in years following stronger relative house price growth in her fixed social network.

However, despite the fact that friends' house price experiences do not vary with individual characteristics on average, one might still be concerned that unobserved shocks to an individual's ability or desire to buy a house in a given year might be correlated with her friends' house price experiences in that year through a channel other than social interactions. Any such alternative story requires a shock to an individual's ability or desire to buy a house in a given local housing market that contemporaneously moves house prices in geographically distant regions where she has friends. For example, many people have friends who work in the same sector. If economic activity in that sector features significant geographic clustering (e.g., tech in Silicon Valley), positive shocks to that sector in a given year might both enable an individual to buy a house and drive up aggregate house prices in those sector-exposed regions where the individual has friends. To rule out this alternative explanation, we show that all results are robust to restricting the sample to individuals who are retired or work in geographically nonclustered professions (e.g., teachers). Our results are also robust to directly including controls for the economic conditions in a person's social network and to interacting our large set of individual demographic controls with year fixed effects, which allows, for example, the effect of different education levels or different occupations on housing market behavior to vary over time.

After ruling out these and other noncausal interpretations of the observed relationship between the house price experiences of an individual's friends and her own housing investment behavior, we explore which channels might explain the observed causal relationship.

We first provide evidence for an important effect of social interactions on an individual's assessment of the attractiveness of property investments, which would naturally affect her housing market investment behavior. To do this, we analyze 1,242 responses to a housing market survey among Los Angeles–based Facebook users. Over half of the survey respondents report that they regularly talk to their friends about investing in the housing market. The survey also asked respondents to assess the attractiveness of property investments in their own zip code. We find a strong positive relationship between the recent house price experiences of a respondent's friends and whether that respondent believes that local property is a good investment. Importantly, this relationship is stronger for individuals who report regularly talking to their friends about investing in property. These results suggest that social interactions provide a link between friends' house price experiences and an individual's own housing market expectations and highlight an important channel through which these experiences can causally affect individuals' housing market investments.

Why would an individual's beliefs about the attractiveness of local housing investments be affected by the house price experiences of her geographically distant friends? While our analysis does not allow us to distinguish between all possible explanations for this behavior, we present some evidence that it is unlikely to be the result of purely rational learning. For example, we show that an individual's investment response to the house price experiences of her friends does not depend on the correlation between her friends' house price experiences and future Los Angeles house price growth. However, there remain a number of possible explanations. For example, our findings could be due to the spread of irrational sentiments as described in Akerlof and Shiller (2010) or due to overconfidence, with individuals overreacting to noisy signals they receive through their social networks (Barberis and Thaler 2003).

We also find no evidence that the causal relationship between the house price experiences of a person's friends and her own housing investment behavior can be explained by a channel other than expectations. First, we document that our results are not driven by individuals investing more in real estate as the value of their expected housing bequest increases with the house price gains of their geographically distant family members. Second, we show that our findings cannot be explained by a story of consumption externalities, such as a desire to "keep up with the Joneses." Finally, we rule out that the observed findings are driven by a desire of individuals to hedge against house price growth in areas they eventually desire to move to.

Overall, our results provide strong evidence for a causal effect of friends' house price experiences on individuals' housing market behavior that works through affecting those individuals' beliefs about the attractiveness of housing investments. In follow-on work, Bailey et al. (2017) show that the shifts in beliefs induced by friends' house price experiences also affect individuals' mortgage leverage choice.

We view our paper as making two contributions. First, we highlight that newly emerging data from online social networking services such as Facebook can overcome the measurement challenges that have held back empirical studies of the economic effects of social networks. In this sense, we add to a recent literature that shows how large data sets collected by online services can help economists understand issues such as households' responses to income shocks (Baker 2018), credit card repayment behavior (Kuchler 2013), the effect of labor mobility on entrepreneurship (Jeffers 2017), housing search behavior (Piazzesi, Schneider, and Stroebel 2017), and online pricing strategies (Einav et al. 2015). In related work, Bailey et al. (2018) aggregate social network data from Facebook to produce a county-level "Social Connectedness Index" that can be shared with other researchers. They use these data to document that other indicators of social and economic activity measured at the regional level, such as trade flows, migration, and patent citations, are related to the degree of social connectedness between regions.

Our second contribution is to use large-scale administrative social network data from Facebook to document that social interactions play an important role in shaping individuals' housing market beliefs and investment behaviors. These empirical findings speak to a number of literatures.

First, we show that differences in friends' house price experiences are an important source of heterogeneity in individuals' housing market expectations. This result contributes to a research effort analyzing how people form expectations about economic outcomes. One popular explanation is that such expectations depend on own experiences. For example, Kuchler and Zafar (2015) show that past local house price changes influence individuals' expectations of future national house price changes. Recent personal experiences also affect expectations in other settings (e.g., Vissing-Jorgensen 2003; Choi et al. 2009; Malmendier and Nagel 2011; Greenwood and Shleifer 2014). We expand on this literature by showing that individuals' expectations are also influenced by the experiences of their friends. These results suggest that differences in social networks can help explain disagreement about asset values among investors. Our findings also provide empirical support for theories in which communication between agents propagates shocks to expectations, in particular in the housing market (e.g., DeMarzo, Vayanos, and Zwiebel 2003; Akerlof and Shiller 2010; Acemoglu et al. 2011; Angeletos and La'O 2013; Shiller 2015; Bayer et al. 2016; Burnside, Eichenbaum, and Rebelo 2016).

Our empirical analysis also documents that individuals with friends who experienced more positive recent house price changes, and who thus believe that housing is a more attractive investment, actually do invest more in real estate and are willing to pay more for a given house. These findings provide support for an important class of models in which expectation heterogeneity influences asset valuations and motivates individuals to trade (e.g., Miller 1977; Harrison and Kreps 1978; Hong and Stein 1999, 2007; Scheinkman and Xiong 2003; Geanakoplos 2009; Simsek 2013). Most directly, our findings provide evidence for a number of papers that focus on the role of heterogeneous expectations and shifts between optimism and pessimism about future house price growth in causing price fluctuations and trading volume in the housing market (e.g., Piazzesi and Schneider 2009; Berger et al. 2017; Kaplan, Mitman, and Violante 2017; Landvoigt 2017; Nathanson and Zwick 2017).

The paper proceeds as follows. Section I argues that data from online social networking services such as Facebook can help researchers measure real-world friendship networks. We also document important features of the Facebook social graph for the United States. Section II describes our empirical approach for identifying a causal effect of a person's friends' house price experiences on her own housing investment behavior. Section III explores the relationship between the average house price experiences in an individual's social network and that individual's housing market investments. Section IV investigates various mechanisms for explaining the observed causal effect. Section V presents conclusions.

I. Measuring Social Networks Using Facebook Data

The key measurement challenge for the empirical literature studying social networks is the difficulty of observing, at a large scale, which individuals are connected to each other. In this section, we show that data from online social networking services such as Facebook and LinkedIn can overcome this measurement challenge and can provide important insights into the structure of social networks. We first discuss the problems with existing approaches to measuring social networks. We then introduce our data on the Facebook social graph and highlight why we believe it provides a realistic representation of real-world friendship networks. Finally, we explore important dimensions of US social networks as described by the Facebook social graph.

A. Approaches to Measuring Social Networks

Traditionally, social scientists have collected data on the structure of real-world social networks using a range of survey techniques (see Morris 2004). There are a number of conceptual and practical challenges with such survey-based approaches to measuring social networks. On the conceptual side, it is well documented that the network structure measured through surveys is sensitive to the exact method of elicitation (e.g., Kogovšek and Ferligoj 2005). The practical challenge is that collecting social network data via surveys is costly to scale. As a result, empirical analyses of real-world social networks have often focused on studying a few publicly available data sets. The most prominent of these is the "Add Health" data from the National Longitudinal Survey of Adolescent Health, which collected information on the social networks among US high school students. Alternatively, researchers have focused on social network data from developing countries, where the cost of collecting information on network structure is less prohibitive (e.g., Alatas et al. 2016; Breza et al. 2017).

More recently, data obtained from online social networks such as Twitter, Google+, and Facebook have provided researchers with opportunities to study the structures of larger-scale social networks (e.g., Ugander et al. 2011; Magno et al. 2012; Shin et al. 2015). While some researchers have worked directly with administrative data from the social networking services, most studies have relied on collecting data by mechanically crawling the social networks' public sites. One problem with data collected through such crawling is that the probability of a particular node being observed depends on the network characteristics of that node. For example, nodes with fewer connections are less likely to be discovered, inducing systematic bias in the observed network structure. These biases highlight the advantage of working directly with administrative data from the social networking services.

More generally, most of the progress in describing the structure of online social networks has been made by researchers in the field of computer science. Much less work has been done by researchers in the social sciences who are as interested in the social and economic implications of network structures as they are in the structures themselves. It is this audience that we have in mind when we describe the Facebook social graph in Section I.C.

B. *The Facebook Social Graph*

Our data contain a de-identified snapshot of all US-based active Facebook users from July 1, 2015. Facebook was created in 2004 as a collegewide online social network for students to maintain a profile and communicate with their friends. It has since grown to become the world's largest online social networking service, with over 1.9 billion monthly active users globally and 234 million monthly active users in the United States and Canada (Facebook 2017). For the users in our data, we observe demographic information, such as their age, education, and county of residence, as well as the set of other Facebook users they are connected to. Using the language adopted by the Facebook community, we call these connections "friends." These data allow us to map out the "social graph" of connections between all US-based Facebook users in our anonymized snapshot.

There are two primary advantages of exploring the Facebook social graph for researchers interested in understanding the economic effects of social networks. The first advantage is Facebook's enormous scale, combined with a user body that is comparatively representative of the US population. Duggan and Page (2016) report that, as of April 2016, more than 68 percent of the US adult population and 79 percent of the US online adult population used Facebook. They also report that, among US online adults, Facebook usage rates were relatively constant across income groups, education groups, racial groups, and urban, suburban, and rural individuals. Usage rates among US online adults were somewhat declining in age, from 88 percent of individuals aged 18–29 years to 62 percent of individuals older than 65 years. This high coverage and relative representativeness of the US population are unique among online social networks. According to Duggan and Page, the three next-largest online social networks, Instagram, Pinterest, and LinkedIn, have at most 40 percent of the US user base that Facebook does; their coverage also drops off much more substantially with age.

The second advantage of Facebook data is that, in the United States, Facebook primarily serves as a platform for real-world friends and acquaintances to interact online (Hampton et al. 2011; Jones et al. 2013). Establishing a friendship link on Facebook requires the consent of both individuals, and there is an upper limit of 5,000 on the number of friends a person can add. Duggan et al. (2015) surveyed Facebook users to characterize their Facebook friendship networks: 93 percent of users said they were Facebook friends with family members other than parents or children; 91 percent said they were Facebook friends with current friends; 87 percent said they were connected to friends from the past, such as high school or college classmates; 58 percent said they were connected to work colleagues; 45 percent and 43 percent said they were Facebook friends with their parents and children, respectively; and 36 percent said they were Facebook friends with their neighbors. Only 39 percent of survey respondents reported to have a Facebook connection to someone they had never met in person. This close correspondence between the Facebook social graph and real-world friendship networks sets it apart from other online social networks, such as LinkedIn, which is more representative of individuals' professional networks, and Twitter, where unidirectional links to individuals that are not real-world acquaintances are common.

C. *Descriptive Statistics on US Social Networks*

In this section, we explore the structure of the Facebook social graph. In particular, we analyze the size and local clustering of individuals' networks as well as patterns of assortativity and homophily. These network characteristics have been described by Jackson, Rogers, and Zenou (2017, 52) as "particularly prominent, fundamental and provid[ing] essential insight" for economists. We also analyze the geographic dispersion of US social networks. We use anonymized data on the full social graph among US-based Facebook users as of July 1, 2015, to construct the individual-level network measures. We then present summary statistics of these network measures across individuals based on a 3 percent random sample of those individuals for whom we observe a full set of demographics such as age, education, and location.

The paper most closely related to this analysis, by Ugander et al. (2011), also explores administrative data on the Facebook social graph. These researchers focused on documenting features of the overall network, without analyzing how characteristics of individuals' positions in the network correlate with individual-level demographics. We argue that these correlations are interesting for a number of reasons. First, they provide useful information for researchers attempting to understand and model the network formation process. Second, policy makers wanting to target information to individuals with particular network positions can use the demographic characteristics of individuals to proxy for their usually unobservable network characteristics. Third, as we highlight in the second half of this paper, understanding the heterogeneity in social network structure across individuals can provide researchers with empirical variation to identify the causal effects of social interactions on economic and financial decision making.

Degree Distribution

An individual's degree centrality, or degree, captures her number of friendship links. The average degree, as well as how it is distributed across individuals, influences how ideas, information, and new technologies spread through a network. All else equal, diffusion is faster in denser networks with more connections. In addition, holding the average degree fixed, an increase in the variance of the degree centrality across individuals is associated with "hub-and-spoke" networks in which a few highly connected nodes play a particularly important role in the diffusion of information.

Jackson (2010) discusses the degree distributions that arise under two prominent models of network formation. In one model, the probability of a link forming between any pair of nodes is equal and independent. This process generates "Poisson random graphs," in which the degree centrality is relatively evenly distributed across nodes. In a second model of network formation, the probability of a given source node forming a connection to a target node is increasing in the degree of the target node. Such a process generates thick-tailed "scale-free" degree distributions in which the frequency of a given degree is proportional to the degree raised to a power.

We begin by analyzing the degree distribution in the Facebook social graph. All reported measures of degree centrality are normalized by the average degree in the sample. Table 1 shows substantial heterogeneity in degree across individuals. At the 10th percentile of the distribution, the degree centrality is 12 percent of its average value, and at the 90th percentile of the distribution, it is 2.23 times as large as the average degree centrality. Panel A of figure 1 plots the log of the degree of a node against the log of the frequency of nodes with that degree in the data. The degree distribution in the Facebook social graph has thicker tails than a Poisson random graph, but high-degree nodes are less common than they would be under a scale-free distribution, which would generate a linear relationship in the log-log space.

TABLE 1
SUMMARY STATISTICS ON US SOCIAL NETWORKS

	Normalized Degree Centrality				Local Clustering Coefficient				Normalized Unique Friends-of-Friends
	Mean	P10	P50	P90	Mean	P10	P50	P90	Mean
Full sample	1.00	.12	.62	2.23	.106	.038	.084	.196	1.00
Age:									
18–34	1.33	.18	.91	2.83	.108	.037	.087	.203	1.35
35–55	.84	.13	.57	1.76	.094	.035	.075	.167	.83
55+	.47	.07	.29	.98	.125	.045	.099	.228	.42
Education:									
No college	.86	.11	.51	1.88	.124	.043	.099	.232	.80
Some college +	1.06	.13	.68	2.35	.100	.036	.079	.184	1.09
County of residence:									
Urban	1.00	.12	.62	2.23	.107	.037	.084	.199	1.01
Rural	1.02	.16	.70	2.19	.125	.050	.106	.220	.81

NOTE.—The table shows summary statistics on the social graph among US-based Facebook users as of July 1, 2015. It contains information on the degree centrality of individuals (normalized by its sample mean), the local clustering coefficient, and the number of unique friends-of-friends (normalized by its sample mean). The full graph is used to construct individual-level statistics, while summary statistics are based on a 3 percent random sample of those individuals for whom we observe a full set of demographics. Summary statistics are presented for the full sample as well as for broad demographic groups.

We also explore which demographic characteristics are associated with an individual's degree. This allows us to consider the roles that various demographic groups play in the diffusion of new ideas or technologies. Table 1 shows that degree centrality is strongly decreasing in age, somewhat increasing in education levels, and relatively constant across individuals living in urban and rural counties. While there is substantial heterogeneity in the average degree across age and education groups, this is not the primary driver of the dispersion of the overall degree distribution: differences in age explain only 8.6 percent of the across-individual variance in degree, while differences in education explain about 1 percent, and differences in urban/rural location explain essentially none of the variance.[1]

Local Clustering

We next explore the extent of local clustering of the friendship networks in our data. The local clustering coefficient of person i measures, across all individuals j and k that are friends with person i, the proportion of

[1] These numbers correspond to the R^2 of separate regressions of degree centrality on dummy variables for each value of age in years, education level ("at most high school," "some college," "some graduate school"), and county of residence.

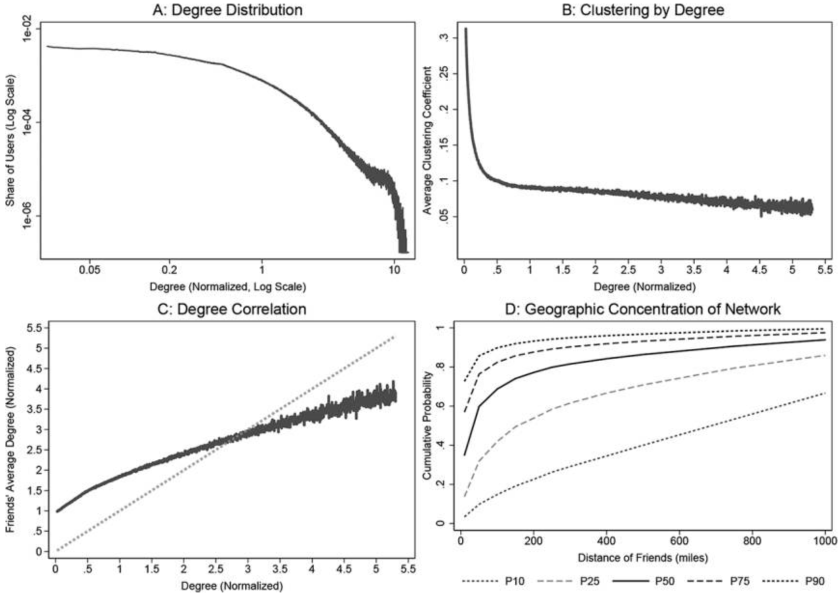

FIG. 1.—Summary statistics on the US social graph. The figure shows summary statistics on the social graph among US-based Facebook users as of July 1, 2015. The full graph is used to construct individual-level statistics, while summary statistics are based on a 3 percent random sample of those individuals for whom we observe a full set of demographics. Panel A shows the correlation between a node's (normalized) degree centrality and the log of the probability of observing a node with that (normalized) degree centrality. Panel B shows the average clustering coefficient for nodes of varying normalized degrees. Panel C shows the average normalized degree centrality of friends by the normalized degree centrality of the own node. Panel D describes the geographic concentration of friendship networks. It shows, for various distances, percentiles of the cumulative distribution of individuals who have at least this many friends living within the respective geographic distance.

pairs that are connected to each other. The local clustering of individuals' networks is of interest to economists for at least two reasons. First, as highlighted by Jackson et al. (2017), having common friends can help sustain social norms and cooperative behavior, for example, because in more clustered networks, news of an individual's undesirable behavior more quickly reaches her friends. Second, clustered networks increase the risk of "persuasion bias" or "correlation neglect" in social learning (DeMarzo et al. 2003; Enke and Zimmermann 2017), whereby individuals fail to account for possible repetitions in the information they receive. Specifically, individuals might ignore the fact that both their own beliefs and the beliefs of their friend might be influenced by a third party to whom they are both connected. This can make people's beliefs and actions excessively sensitive to signals they receive through their social network.

Table 1 shows substantial heterogeneity in the extent of local network clustering across individuals. For the average person, the clustering coefficient is 0.106 (suggesting that 10.6 percent of friend-pairs are themselves friends), but this ranges from 0.038 at the 10th percentile of the distribution to 0.196 at the 90th percentile. The clustering coefficient is U-shaped in age, with younger and older individuals having more clustered networks than middle-aged individuals, despite the fact that younger people have the largest networks on average. Networks of rural individuals and individuals without any college attendance are somewhat more clustered. On average, larger networks are less clustered (see panel B of fig. 1), suggesting that friends of higher-degree nodes are less likely to be friends with each other. Indeed, network size explains 22.6 percent of the across-individual variation in local clustering, while age explains 7.4 percent, education explains 2.0 percent, and urban/rural location explains 0.2 percent.

More clustered networks mean that, for the same network size, individuals are exposed to fewer unique friends-of-friends and therefore to fewer ideas and opportunities that might travel over several links. For example, table 1 shows that urban individuals are connected to more unique friends-of-friends than rural individuals, despite the fact that they have slightly smaller networks on average. This is in part explained by the lower clustering of urban networks.

Assortativity and Mixing Patterns

An important aspect of social networks is the extent to which individuals are friends with others who are similar to them along dimensions such as the position in the network or demographic characteristics.

We first analyze such assortativity based on network characteristics. We focus on the extent of "degree correlation," which captures the tendency of high-degree nodes to be connected to other high-degree nodes. This network feature is important, since diffusion processes are usually faster in networks with significant degree correlation (Barabási 2016). Panel C of figure 1 plots the relationship between the degree of individuals and the average degree of their friends. We observe significant positive degree correlation: across the sample, the correlation of an individual's own degree and the average degree of her friends is 65 percent. While high-degree nodes are generally connected to other high-degree nodes, most individuals have friends who, on average, have more links than they do. This is a manifestation of Feld's (1991) paradox that "your friends have more friends than you do." Indeed, until individuals have substantially more than 2.5 times the average degree (which is about the 95th percentile of the distribution), their average friend has more friends than they do.

We also explore the extent to which individuals' friendship networks disproportionately include other individuals who are similar in terms of demographic characteristics. McPherson, Smith-Lovin, and Cook (2001) document that this type of "homophily" is a common feature across many social networks, and it affects the extent to which individuals are exposed to a diverse set of views through their friends. We find substantial assortativity based on age: 78.9 percent of the friends of individuals aged between 18 and 34 are themselves between those ages (see table 2). Only 4.7 percent are above 55 years old. When we focus on individuals aged above 55 years, we find the reverse pattern: 46.8 percent of their friends are older than 55 years, while only 17.6 percent are younger than 35 years. Similar homophily can be detected across education and rural/urban residents, with evidence that individuals are, on average, more likely to be friends with others who are similar to them on demographic characteristics.

The Geographic Dimension of Social Networks

The last dimension of the social graph that we explore is the geographic concentration of friendship links. The extent to which social networks

TABLE 2
Homophily in US Social Networks

	Share of Friends by Age Group (%)			Share of Friends by Education Group (%)		Share of Friends by County of Residence (%)	
	18–34	35–55	55+	No College	Some College +	Urban	Rural
Full sample	48.9	34.9	16.2	27.1	72.9	93.5	6.5
Age:							
18–34	78.9	16.3	4.7	28.1	71.9	94.0	6.0
35–55	25.9	58.5	15.6	26.4	73.6	93.3	6.7
55+	17.6	35.6	46.8	26.1	73.9	92.5	7.5
Education:							
No college	49.1	34.4	16.3	35.9	64.1	92.1	7.9
Some college	48.9	35.1	16.1	23.3	76.7	94.1	5.9
County of residence:							
Urban	49.2	34.8	16.0	26.7	73.3	96.0	4.0
Rural	45.7	35.9	18.3	33.1	66.9	56.2	43.8

Note.—The table shows summary statistics on the social graph among US-based Facebook users as of July 1, 2015. It contains information on the share of friends that belong to different broad demographic groups (among all friends for whom we have information on that demographic characteristic). The full graph is used to construct individual-level statistics, while summary statistics are based on a 3 percent random sample of those individuals for whom we observe a full set of demographics. Summary statistics are presented for the full sample as well as for broad demographic groups.

are geographically concentrated affects a number of important outcomes, such as whether social networks can provide insurance against regional shocks.

Table 3 shows that the average individual has friends in over 70 different counties. Despite this relatively large number of counties with at least one friend, the average individual has 34.7 percent of friends living in the same county and 63.3 percent of friends living in the same state. When measured in geographic distances, the average person has 52.7 percent of friends living within 50 miles and 67.4 percent of friends living within 200 miles. There is substantial heterogeneity in these numbers: panel D of figure 1 shows percentiles of the cumulative distribution of friends living at ranges up to 1,000 miles. Some of this heterogeneity is driven by differences in the geographic concentration of friendship networks across demographic groups, with older and more educated users having less geographically concentrated networks. For example, the share of friends living within 200 miles falls from 69.6 percent for individuals aged 18–34 years to 62.6 percent for individuals aged 55 years and older. While urban individuals have slightly more of their friends living within 50 miles than rural individuals (52.8 percent vs. 51.7 percent), they have fewer friends living within 200 miles (67.1 percent vs. 72.8 percent). Overall, age and education each explain about 2.5 percent of the across-individual variation in the share of friends living within 200 miles, while rural/urban location explains about 0.3 percent.

TABLE 3
Geographic Distribution of US Social Networks

		Share of Friends Living Within (%)				
	Number of Counties	Own County	Own State	50 Miles	200 Miles	500 Miles
Full sample	70.5	34.7	63.3	52.7	67.4	77.4
Age:						
18–34	81.6	37.1	65.8	54.3	69.6	79.4
35–55	67.3	33.8	62.6	53.1	67.1	76.9
55+	48.7	30.4	58.4	47.9	62.6	73.2
Education:						
No college	57.6	39.4	68.0	58.3	71.9	80.7
Some college	76.2	32.7	61.3	50.3	65.5	76.0
County of residence:						
Urban	70.2	35.0	63.0	52.8	67.1	77.0
Rural	75.1	30.1	68.0	51.7	72.8	83.4

Note.—The table shows summary statistics on the social graph among US-based Facebook users as of July 1, 2015. It contains information on the geographic distribution of friendship networks. The full graph is used to construct individual-level statistics, while summary statistics are based on a 3 percent random sample of those individuals for whom we observe a full set of demographics. Summary statistics are presented for the full sample as well as for broad demographic groups.

Table 4 explores how the geographic concentration of social networks varies across different regions in the United States. There is substantial heterogeneity across counties in the share of friends who live nearby. At the 5th percentile of the distribution, the median person in a county has 31.2 percent of friends living within 50 miles; at the 95th percentile, this number is 72.6 percent. There is similar heterogeneity across counties in the share of friends living within 200 miles, with a 5–95 percentile range of 48.5–86.5 percent. Panel A of figure 2 plots the share of friends living within 200 miles for the median person living in each county in the continental United States. Social networks are most geographically concentrated in the South, the Midwest, and Appalachia. In fact, the 12 counties with the most concentrated networks are all in Kentucky. On the other hand, social networks in the sparsely populated parts of the noncoastal western United States are the least geographically concentrated. The exception is Utah, which has fairly geographically concentrated social networks.

We next analyze how the geographic concentration of the social networks of a county's population correlates with county-level demographics. Panels B and C of figure 2 show county-level binned scatter plots of the relationship between the share of friends of the median person in a county who lives within 200 miles and two county-level demographic mea-

TABLE 4
Geographic Distribution of US Social Networks: County-Level Heterogeneity

	Share of Friends Living within 50 Miles (%)			Share of Friends Living within 200 Miles (%)		
	Median Person	95–5 Range	75–25 Range	Median Person	95–5 Range	75–25 Range
Mean	55.4	73.5	34.5	72.5	70.7	28.0
P5	31.2	62.4	22.9	48.5	51.2	13.9
P10	37.9	65.7	25.3	56.0	55.4	15.1
P25	46.3	70.9	30.1	68.0	63.0	18.4
P50	57.7	74.2	33.9	75.1	73.7	27.1
P75	65.4	77.2	39.4	80.7	78.5	37.3
P90	69.9	80.7	43.1	84.5	81.3	43.6
P95	72.6	81.8	45.5	86.5	82.7	46.6

Note.—The table shows summary statistics on the social graph among US-based Facebook users as of July 1, 2015. It contains information on how the geographic concentration of friendship networks varies across and within US counties. Columns 1 and 4 show how the geographic concentration of the social networks of the median person in each US county varies across counties. For the statistics in the other columns, we first calculate, for every county, the 95–5 percentile range (cols. 2 and 5) and the 75–25 percentile range (cols. 3 and 6) of social network concentration across the county's population and then show distributions across counties. The full graph is used to construct individual-level statistics, while summary statistics are based on a 3 percent random sample of those individuals for whom we observe a full set of demographics. County-level summary statistics are constructed by population-weighting the individual counties.

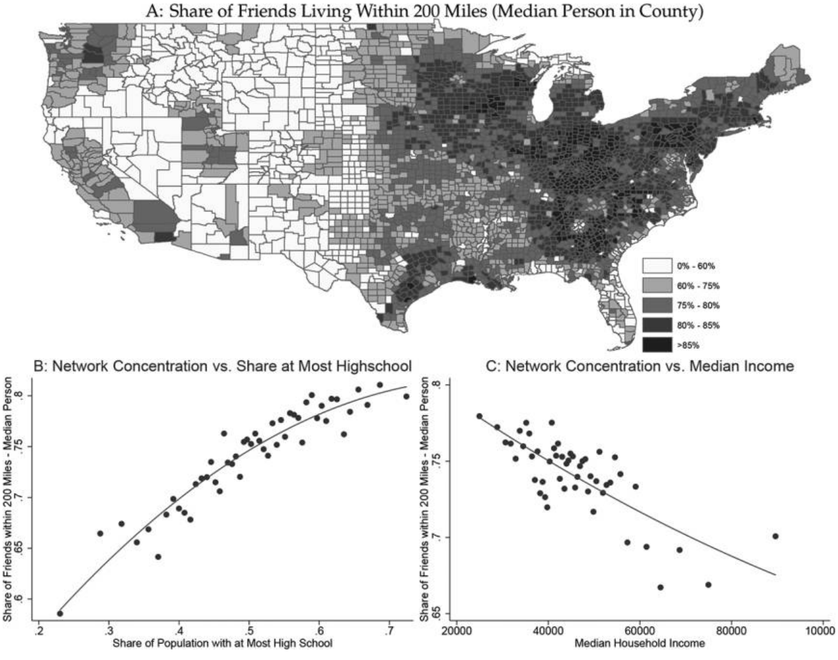

FIG. 2.—Geographic concentration of US social networks—county level. The figure shows summary statistics on the social graph among US-based Facebook users as of July 1, 2015. The full graph is used to construct node-level statistics, while summary statistics are based on a 3 percent random sample of those nodes for which we observe a full set of demographics. Panel A shows a heat map of the share of friends living within 200 miles for the median person living in each county. Panels B and C show county-level binned scatter plots (50 bins) of the relationship between the share of friends of the median person living within 200 miles on the vertical axes and the share of population with at most a high school diploma and the median household income in US dollars, respectively, on the horizontal axes. The demographic measures come from the 2010 5-year estimates of the American Community Survey.

sures from the 2010 5-year wave of the American Community Survey: the share of individuals with at most a high school diploma and the median household income. Consistent with the individual-level patterns, counties with higher education levels and higher incomes have less geographically concentrated friendship networks.

While there is significant across-county variation in the geographic structure of social networks, there remains substantial within-county variation. Table 4 shows the across-county distribution of the within-county interquartile range of the share of friends living within 200 miles. For the average county, this interquartile range is 28 percent. Even for counties with relatively homogeneous social networks, there is substantial variation in the geographic dispersion of the residents' social networks.

Key Takeaways

In this section, we analyzed a number of important characteristics of the friendship networks of US Facebook users. While there are systematic patterns in how the network structure varies with individual demographic characteristics, substantial heterogeneity in network characteristics remains within demographic groups. As we show next, this heterogeneity can provide researchers with interesting variation to explore the economic effects of social networks.

II. Social Networks and Housing Markets: Research Design and Data

In the previous section, we argued that the increasing availability of data from online social networking services substantially expands the potential for empirical research on the economic effects of social networks. In the remainder of the paper, we substantiate this point by using data from Facebook to explore the role of social interactions in influencing individuals' housing investment decisions. We begin by describing our empirical strategy for isolating a causal effect of the house price experiences of an individual's friends on her housing market investments.

Our baseline specifications are regressions of individual i's housing investment decisions at time t_2 on measures of the average house price experience within that individual's social network between t_1 and t_2, $\text{FriendHPExp}^N_{i,t_1,t_2}$.[2] We control for individual i's demographics as well as location × time fixed effects, represented by $\mathbf{X}_{i,t_2}$. This allows us to isolate the effects of friends' house price experiences on the housing investment decisions of otherwise similar individuals at the same point in time:

$$\text{HousingInvestment}_{i,t_2} = \beta \text{FriendHPExp}^N_{i,t_1,t_2} + \gamma \mathbf{X}_{i,t_2} + \epsilon_{i,t_2}. \quad (1)$$

To measure the house price experiences of a person's friends, we combine data on the county of residence of her friends with county-level house price indices from Zillow. Let $\text{ShareFriends}_{i,N,c}$ be the share of person i's Facebook friends in network N who live in county c. Similarly, let $\Delta\text{HP}_{c,t_1,t_2}$ capture the house price changes in county c between t_1 and t_2. We then construct our primary explanatory variable as

$$\text{FriendHPExp}^N_{i,t_1,t_2} = \sum_c \text{ShareFriends}_{i,N,c} \times \Delta\text{HP}_{c,t_1,t_2}. \quad (2)$$

[2] In this paper, we focus on the effects of the average house price experiences of a person's friends. However, other moments of the distribution of friends' house price experiences might also matter. Indeed, Bailey et al. (2017) show that individuals' house price expectations and mortgage leverage choices are affected by both the first and the second moments of the distribution of house price experiences across their friends.

This measure of friends' house price experiences can be constructed for different networks N of individual i. The broadest such network includes all of her Facebook friends, but other subnetworks might include, for example, her out-of-commuting-zone friends or her work friends.[3]

In order to interpret estimates of β in regression (1) as the causal effect of friends' house price experiences, we rule out potential alternative, noncausal channels that might also induce a correlation between a person's housing market investments and her friends' house price experiences.

A first concern is that FriendHPExp might be correlated with an individual's own house price experiences or her own past capital gains, both of which could directly affect her housing investment decisions. In particular, since most people have many local friends, shocks to local house prices will shift FriendHPExp, with larger shifts for people with a larger share of local friends. Therefore, any confounding effect of past local house price movements on housing investments that is stronger for people with a larger share of local friends would affect our interpretation of β. For example, suppose that people who have lived in Los Angeles for longer both have more friends in Los Angeles and are more likely to extrapolate from Los Angeles house prices when forming their expectations about future house price growth. This could induce a correlation between a person's housing market investments and FriendHPExp that is not due to social interactions. Similarly, imagine that a person who has lived in Los Angeles for longer is more likely to already own a house there. In that case, higher Los Angeles house price growth can have a stronger effect on this person's housing market investments both because her larger local network has experienced bigger house price increases and because she has larger past capital gains on her existing home. If we cannot control for such past capital gains, we would erroneously attribute all observed effects to social interactions.

To address this challenge, we estimate regression (1) using an instrumental variables (IV) strategy, where we instrument for the house price experiences of all of a person's friends with the house price experiences of only her geographically distant friends. In the baseline specifications, we use the house price experiences of her out-of-commuting-zone friends

[3] Our measure of friends' house price experiences treats each friendship link in a given network N identically. Weighting different friends by their tie strength does not systematically affect our results since the geographic distributions of strong and weak ties are usually similar. Since we observe only one snapshot of the Facebook social graph, we cannot exploit time-series variation in an individual's social network. Thus $\text{FriendHPExp}_{i,t_1,t_2}$ measures the house price experiences between t_1 and t_2 of person i's social network as of the date of the snapshot, July 1, 2015. The interpretation of our empirical estimates thus requires that the counties that an individual was exposed to through her friends in 2015 provide an unbiased estimate of the counties that she was exposed to at the time we measure her housing investment behavior.

as the instrument, but we also show robustness to using her out-of-state friends' experiences. The first and second stages of this IV regression, respectively, are given by

$$\text{FriendHPExp}^{\text{All}}_{i,t_1,t_2} = \beta^{\text{FS}}\text{FriendHPExp}^{\text{OutCZ}}_{i,t_1,t_2} + \delta\mathbf{X}_{i,t_2} + \varepsilon_{i,t_2} \tag{3}$$

and

$$\text{HousingInvestment}_{i,t_2} = \beta^{\text{IV}}\,\widehat{\text{FriendHPExp}}^{\text{All}}_{i,t_1,t_2} + \gamma\mathbf{X}_{i,t_2} + \epsilon_{i,t_2}. \tag{4}$$

The instrument has F-statistics above 1,500 across all first-stage regressions. The reason is that the construction of the instrumented variable directly builds on the instrument

$$\begin{aligned}\text{FriendHPExp}^{\text{All}}_{i,t_1,t_2} = {} & \text{ShareFriendsCZ}_i \times \Delta\text{HP}_{\text{CZ},t_1,t_2} \\ & + (1 - \text{ShareFriendsCZ}_i) \times \text{FriendHPExp}^{\text{OutCZ}}_{i,t_1,t_2}.\end{aligned}$$

Indeed, if all people had the same share of local friends, the first-stage regression with county × time fixed effects would have an R^2 of 1. The second-stage regression includes a predicted $\widehat{\text{FriendHPExp}}^{\text{All}}_{i,t_1,t_2}$ that can be thought of as generated under the assumption that all people have the same share of local friends. Our estimates of β^{IV} are therefore identified only by variation in $\text{FriendHPExp}^{\text{All}}_{i,t_1,t_2}$ that is independent of individual-specific variation in the share of local friends. We can thus rule out concerns that our estimates are confounded by any channel that would induce individuals with more local friends to react more to past local house price changes for reasons other than social interactions.[4]

Even with this IV research design, a further concern relates to people who recently moved to their current commuting zone from geographically distant parts of the country where they have many friends. For these people, there might still be a strong correlation between their own house price experiences and capital gains and the house price experience of their friends who live outside of their current commuting zone. To rule

[4] We choose to interpret estimates from the second-stage IV regression (4) rather than from the reduced-form regression that directly includes $\text{FriendHPExp}^{\text{OutCZ}}_{i,t_1,t_2}$ in regression (1). The reason is that we find the interpretation of the magnitude of the IV estimates to be more natural. In the reduced-form specification, the magnitude of β will be similar to the magnitude of β^{IV} scaled by the average share of out-of-commuting-zone friends. These reduced-form estimates would capture the average effect of the house price experiences of only the out-of-commuting-zone friends on the outcome of interest. One assumption in our interpretation of β^{IV} is that the effect of friends' house price experiences on own housing investments, through social interactions, is similar for geographically close and distant friends. There is some evidence that this is indeed a valid assumption, since our results do not depend on whether we use out-of-commuting-zone friends' or out-of-state friends' experiences as the instrument. However, if one were instead to expect a larger reaction to the experiences of geographically close friends, the magnitude of β^{IV} would understate the effect of the response to the house price experiences of all friends.

out such concerns, we verify that our results are robust to excluding recent movers from our regressions.

We also consider whether the nonrandom variation in individuals' geographically distant social networks documented in Section I.C poses a challenge to our causal interpretation of β in regression (4). A first important observation is that our identification does not require that individuals' social networks do not systematically vary with those individuals' observed and unobserved characteristics. For example, it is not necessarily a problem that people with graduate degrees are more likely to have friends in Boston and are more likely to buy a house. The reason is that our dependent variable is driven by where in the United States people have friends interacted with how house prices in these areas change in a given year. Since house price growth in Boston is sometimes above and sometimes below the US average, the same individual's social network will sometimes experience above-average and sometimes below-average house price changes. By comparing the housing investment behavior of individuals with friends in Boston across different years, we can thus remove the effect of any time-invariant individual-level determinants of housing investments that are correlated with having friends in Boston. Indeed, we document below that the variation in the average house price experiences across different individuals' friends is unrelated to observed or unobserved fixed characteristics of those individuals. Consistent with this, our estimates are unchanged in those empirical specifications in which we include individual fixed effects and thus exploit only within-individual across-time variation in friends' house price experiences.

A second, more subtle concern with our causal interpretation is that shocks to a person's desire or ability to buy a house in a given year might vary systematically with the house price movements in that year in those geographically distant areas where this person has friends. This challenge is weaker than that faced by the peer effects literature, which has to address concerns about common unobserved shocks to individuals and their friends. For example, in our setting it is not problematic that people and their friends have children around similar times and therefore also buy houses around similar times. The reason is that FriendHPExp does not depend on the housing market decisions of an individual's friends. Instead, it is driven only by the house price changes in the counties where those friends live. Therefore, challenges to our identification have to come from shocks that not only affect an individual's own housing market decisions but also move equilibrium house prices in geographically distant counties where that individual has friends.

We were able to identify one such potential challenge to our interpretation coming from individuals working in professions or industries that feature significant geographic clustering. Suppose that people who work in the tech sector have more friends in Silicon Valley. During tech booms, tech employees in Los Angeles might have more resources to buy a house, and the increase in housing demand by the many tech employees in Silicon Valley drives up house prices there. Without controlling for year × tech sector fixed effects, one might falsely attribute large housing investments by Los Angeles–based tech employees in those years to social interactions. We address this challenge using three complementary strategies. First, we estimate specifications that include year-specific controls for a rich set of observable individual characteristics. These interacted controls have no effect on our estimates of β, suggesting that year-specific shocks to different demographic groups that correlate with house price changes in their geographically distant social networks are not driving our results. Second, we show that our results are robust to focusing on the sample of individuals who are retired or work in geographically nonclustered professions (e.g., teachers and legal professionals). Third, to further address concerns about possible confounding effects from income shocks to connected counties, we present specifications that control for friend-weighted income changes over the past 24 months, as measured by changes in the gross income per capita from the Internal Revenue Service Tax Statistics of Income. We show that this additional control does not significantly affect our estimated response of housing investment behavior to friends' house price changes. Jointly, these robustness checks suggest that our estimates are not driven by changes to the economic conditions of an individual's friends, which may correlate with both this individual's own behavior and her friends' house price experiences.

III. Social Networks and Housing Markets: Evidence

We next use the empirical strategy described above to show that the house price experiences within an individual's social networks have a causal effect on her housing investment decisions. We first document an effect of friends' house price experiences on the extensive margin decision to be a homeowner or a renter. We also show that the intensive margin of an individual's housing investment—the square footage of the home bought—as well as the transaction price are positively affected by higher house price experiences in her social network.

A. *Social Networks and Housing Markets: Extensive Margin Analysis*

Housing Data

To measure housing investment activity at the individual level, we introduce data from two snapshots of Acxiom InfoBase, one from 2010 and one from 2012. These data are maintained by Acxiom, a marketing services company, and contain a range of individual-level information compiled from a large number of sources (e.g., public records, surveys, and warranty registrations). The data include details on demographics (e.g., age, marital status, education, occupation, income), household size, and home ownership status. For current homeowners, the data also contain information from public deeds records on the housing transaction that led to the current home ownership spell (e.g., transaction date and price), as well as property details from public assessor records (e.g., property and lot size).

Sample Description and Summary Statistics

We merge the Facebook and Acxiom data through a unique, anonymized link based on common characteristics in both data sets.[5] Since the housing transaction deeds are originally recorded at the county level, we focus our empirical analysis on understanding the housing market behavior of the residents of Los Angeles County, the largest US county by population. This ensures that our analysis is not affected by inconsistent recording of data across counties. Our final sample consists of an anonymized panel of about 1.4 million Facebook users who lived in Los Angeles County in 2010 and whom we can match across the 2010 and 2012 Acxiom snapshots.[6] Below, we exploit the panel structure of this data set to analyze how the 2010–12 transition probability between renting and owning is affected by the individuals' friends' house price experiences between 2008 and 2010. We therefore call this sample the "change-of-tenure sample."

Table 5 contains summary statistics on this sample; additional summary statistics are provided in the appendix (available online). In 2010, the average person was 41 years old and had a household income of almost $70,000. About 29.5 percent of the individuals were renters in 2010; by 2012, 17.8 percent of these 2010 renters had bought a home. Of the 70.5 percent of people who owned their home in 2010, 93.5 percent continued to own their home in 2012. The average person has 304 US-based friends. This number of friends ranges from 35 to 943 between the 5th and the 95th percentile of the distribution.

Figure 3 shows a heat map of the geographic distribution of the aggregated social networks of all individuals in the sample. Consistent with the

[5] Linking the housing data to the friendship network involved a scrambled merge-key based on common characteristics. Fifty-three percent of merges relied on email address. Other characteristics were full date of birth (51 percent) or year-month of date of birth (28 percent), last name (45 percent) and first name (84 percent), location at the level of zip code (44 percent), county (37 percent), core-based statistical area (8 percent), and telephone number (2 percent). Most matches are based on multiple characteristics.

[6] We drop the 17 percent of individuals with fewer than 10 out-of-commuting-zone friends, for whom the measure of friends' geographically distant house price experience is noisy; however, our results are robust to variation in this cutoff.

TABLE 5
Summary Statistics: Change-of-Tenure Sample

	Mean	Standard Deviation	P5	P25	P50	P75	P95
Number of friends	304	406	35	90	184	358	943
Number of counties with friends	55.5	59.9	13	22	37	67	151
Share of friends living within (%):							
Los Angeles commuting zone	62.9	19.8	22.4	51.4	67.9	78.2	87.1
California	70.4	19.2	28.7	61.8	76.4	84.3	91.0
200 miles	65.5	19.9	23.9	54.6	70.9	80.5	88.6
500 miles	74.7	19.3	32.2	66.7	81.1	88.6	94.5
1,000 miles	79.1	18.2	38.0	73.1	85.7	91.6	96.2
Share of out-of-commuting-zone friends by census division (%):							
Pacific	32.4	18.6	5.7	17.8	30.9	44.7	66.7
Mountain	20.1	14.8	2.5	8.5	17.0	28.6	48.5
West North Central	3.4	6.6	.0	.0	1.8	4.1	11.8
East North Central	7.3	10.3	.0	1.9	4.6	8.3	25.0
Mid-Atlantic	9.2	12.1	.0	1.5	5.5	11.8	34.3
New England	2.8	6.0	.0	.0	1.1	3.4	10.0
West South Central	10.3	10.6	.0	4.0	7.6	13.2	29.2
East South Central	2.3	4.7	.0	.0	1.0	3.0	8.3
South Atlantic	12.0	11.2	.0	5.3	9.3	15.4	32.3
Δ friend house prices: 2008–10 (%):							
All friends	−7.1	1.8	−10.1	−7.7	−6.8	−6.1	−5.2
Out-of-commuting-zone friends	−10.3	3.4	−16.3	−12.3	−10.1	−8.1	−5.2
Out-of-state friends	−11.5	4.1	−18.7	−14.0	−11.2	−8.9	−5.2
Δ friend house prices: 2010–12 (%):							
All friends	4.3	1.4	2.1	3.9	4.4	4.9	6.1
Out-of-commuting-zone friends	4.6	2.4	.7	3.3	4.6	5.8	8.1
Out-of-state friends	4.0	2.7	−.2	2.5	4.0	5.4	8.1
Income 2010 ($1,000s)	69.9	41.5	10	35	63	88	150
Income change 2010–12 ($1,000s)	.71	23.1	−35	0	0	0	38
Household size 2010	3.02	1.74	1	2	3	4	6
Household size change 2010–12	−.10	1.26	−2	0	0	0	2
Age 2010	41.0	15.1	20	31	41	51	66
Home ownership development, 2010–12:							
Stayed renter	.24	.43	0	0	0	0	1
Became homeowner	.05	.22	0	0	0	0	1
Stayed homeowner	.66	.47	0	0	1	1	1
Became renter	.05	.21	0	0	0	0	0
Family structure development, 2010–12:							
Stayed single	.42	.49	0	0	0	1	1
Got married	.06	.24	0	0	0	0	1
Stayed married	.47	.50	0	0	0	1	1
Got divorced	.06	.23	0	0	0	0	1
Education 2010:							
Has high school	.47	.50	0	0	0	1	1
Has college degree	.37	.48	0	0	0	1	1
Has graduate degree	.15	.36	0	0	0	0	1

Note.—The table shows summary statistics for the change-of-tenure sample, which consists of Facebook users who lived in Los Angeles County in 2010 and whom we can match across the 2010 and 2012 Acxiom snapshots. $N = 1{,}469{,}359$. For each characteristic, we show the mean, standard deviation, and percentiles of the distribution.

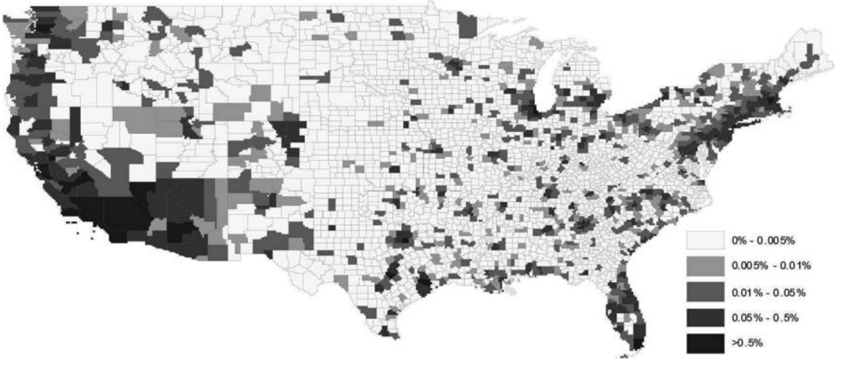

FIG. 3.—Share of friendship links of Los Angeles residents. The figure shows the absolute share of US-based friends of individuals in the change-of-tenure sample who live in each county. The change-of-tenure sample consists of Facebook users who lived in Los Angeles County in 2010 and whom we can match across the 2010 and 2012 Acxiom snapshots.

findings from the US-wide social graph explored in Section I.C, a significant fraction of friendship links are to geographically close individuals. Indeed, table 5 shows that the average person in our sample has 62.9 percent of her friends within the Los Angeles commuting zone and 65.5 percent of her friends living within 200 miles. Despite this relative clustering of friends near Los Angeles, the average person has friends in more than 55 different US counties. There is substantial heterogeneity in where different individuals have these friends. For example, panels A, B, and C of figure 4 map the social networks for three different individuals in our sample whose out-of-commuting-zone friends are clustered around Minnesota, North Carolina, and Utah, respectively. Similarly, table 5 shows that while the average person in our sample has 32.4 percent of her out-of-commuting-zone US friends living in the Pacific census division (comprising Alaska, California, Hawaii, Oregon, and Washington), this number ranges from 5.7 percent to 66.7 percent between the 5th and the 95th percentiles of the distribution. It is this across-individual heterogeneity in the location of geographically distant friends, combined with differences in house price movements across the United States, that is the key driver of variation in friends' house price experiences.

Indeed, while the average person in our sample has friends who experienced a 7.1 percent house price decline between December 2008 and December 2010, this number ranges between −10.1 percent and −5.2 percent from the 5th to the 95th percentiles of the distribution. The 5–95 percentile range of out-of-commuting-zone friends' house price experiences is even larger, ranging from −16.3 percent to −5.2 percent. Panels A and B of figure 5 plot the full distribution of friends' house price experiences separately for all friends and out-of-commuting-zone friends, respectively.

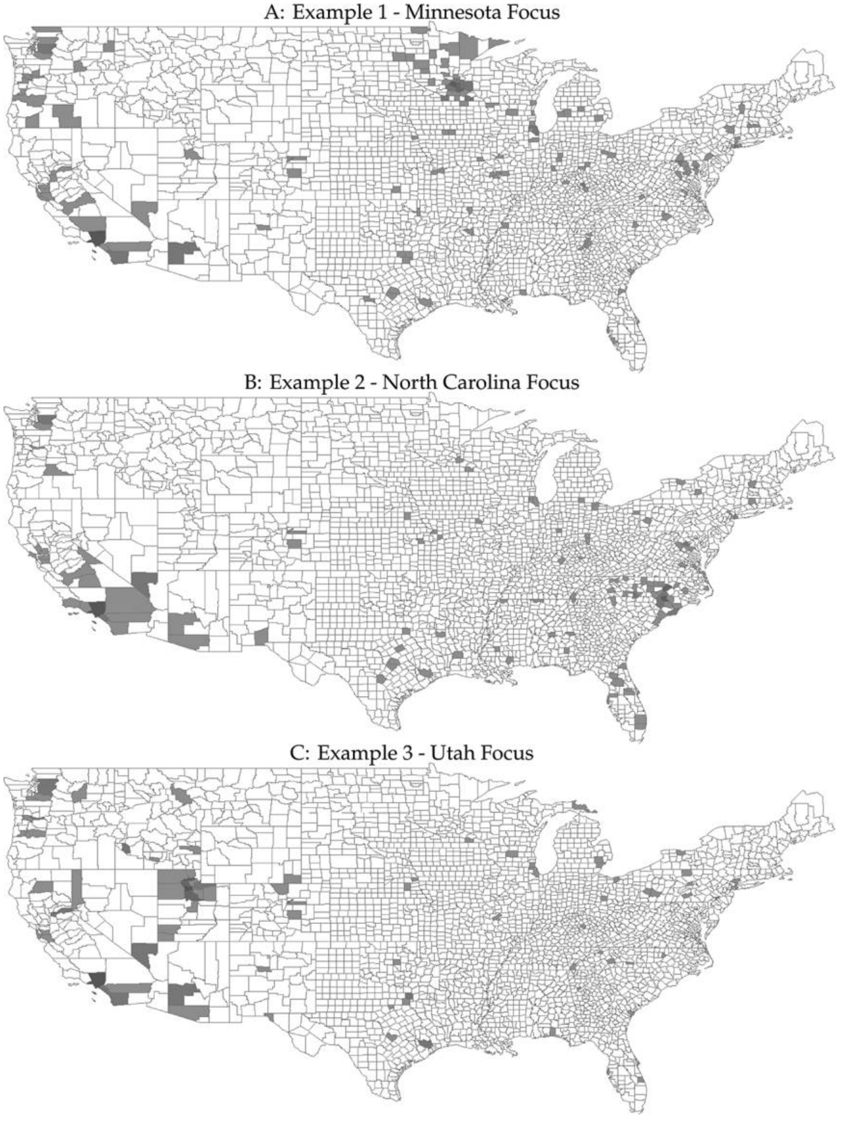

FIG. 4.—Examples of individual-level friend distributions. The figure shows the geographic distribution of the friends of three Facebook users living as renters in Los Angeles County in 2010. Panel A shows an individual with disproportionately many friends clustered in Minnesota. Panel B shows an individual with disproportionately many friends clustered in North Carolina. Panel C shows an individual with disproportionately many friends clustered in Utah.

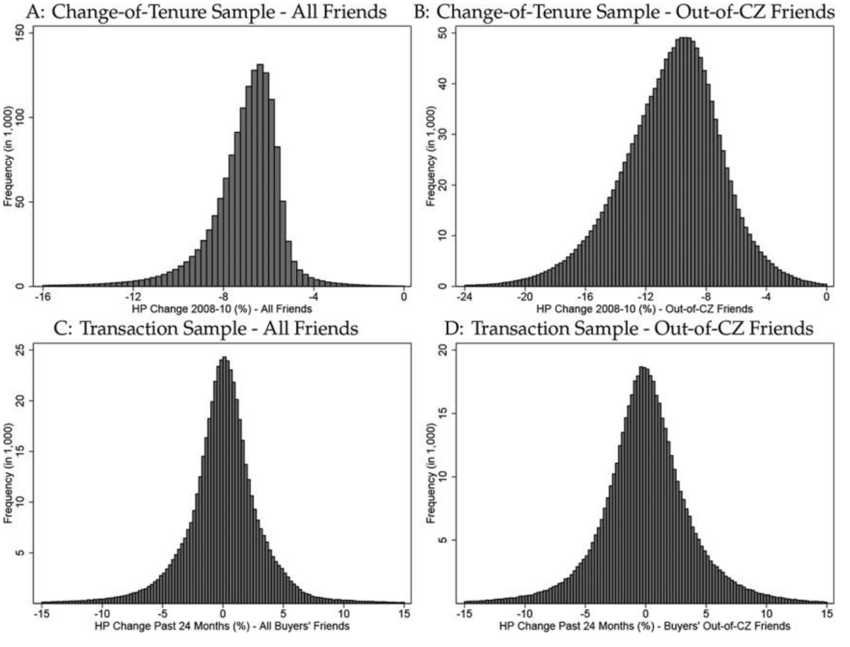

FIG. 5.—Distribution of friends' house price experiences. Panels A and B show the distribution of the average house price experiences between 2008 and 2010 of the friends of individuals in the change-of-tenure sample, which consists of Facebook users who lived in Los Angeles County in 2010 and whom we can match across the 2010 and 2012 Acxiom snapshots. Panel A focuses on the experiences of all friends. Panel B focuses on the experiences of out-of-commuting-zone friends. Panels C and D show the distribution of the average house price experiences of the friends of the buyers in the transaction sample in the 24 months prior to the transaction. The transaction sample consists of all housing transactions by Facebook users in Los Angeles County between 1993 and 2012 that led to a home ownership spell that was still ongoing as of the 2010 or 2012 Acxiom snapshots. Since this pools across transactions in different years, all friend experiences are shown conditional on a quarter-of-transaction fixed effect. Panel C focuses on the experiences of all friends and panel D on the experiences of out-of-commuting-zone friends. The bucket size in all panels is 0.25 percentage points.

Our empirical approach exploits this variation in friends' house price experiences to document a causal role of social interactions in shaping housing investment behavior. A key assumption behind our causal interpretation is that no individual's social network consistently experiences an above-average or below-average house price appreciation. To verify this, we calculate, for each individual in our sample and each year between 1993 and 2012, the house price experiences in that individual's social network over the previous 12 months. We then regress these individual-year observations on individual fixed effects. This regression yields an R^2 of about 0.1 percent. This confirms that, on average, the variation in the house price experiences across different individuals' friends is unrelated to observed or unobserved fixed characteristics of those individuals.

Extensive Margin Results

We first focus on the Los Angeles–based renters in 2010 in the change-of-tenure sample. Regression (5) considers whether their propensity to become a homeowner by 2012 is affected by the house price experiences of their friends between 2008 and 2010:

$$\mathbf{1}_{\text{Owner}_{i,2012}} = \alpha + \beta \text{FriendHPExp}^{\text{All}}_{i,2008,2010} + \gamma \mathbf{X}_{i,2010} + \omega \Delta \mathbf{X}_{i,2010,2012} + \psi_{\text{zip}_{2010},\text{zip}_{2012}} + \epsilon_i. \quad (5)$$

The dependent variable is an indicator of whether individual i is a homeowner in 2012. We control for paired 2010 × 2012 zip code fixed effects (e.g., an indicator variable for all individuals who lived in zip code 90001 in 2010 and in zip code 90005 in 2012), which allows us to isolate the decision of where to live from the decision of whether to buy a house. We also control for the 2010 demographics of individual i, $\mathbf{X}_{i,2010}$, and changes in these demographics between 2010 and 2012, $\Delta\mathbf{X}_{i,2010,2012}$. Our controls also include information on the size of the individuals' Facebook networks, such as the number of friends, the number of out-of-commuting-zone friends, and the number of counties in which they have at least one friend. As described in Section II, to help us isolate the causal effect of friends' house price experiences, we use the house price experiences of friends who live outside the Los Angeles commuting zone to instrument for the house price experiences of all friends.

Panel A of table 6 shows results from regression (5).[7] The estimate of β in column 1 suggests that every percentage point increase in the house price experiences of an individual's friends increases her probability of becoming a homeowner by 2012 by about 0.6 percentage points, relative to a baseline transition probability of 17.8 percent. A one standard deviation increase in the house price appreciation experienced by a person's friends between 2008 and 2010 thus increases the probability of buying a home over the next 2 years by 1.2 percentage points. This is about 18 percent of the magnitude of the effect of adding a family member.

While the $\psi_{\text{zip}_{2010},\text{zip}_{2012}}$ fixed effects help us to separate the choice of location from the choice of owning or renting, in column 2 we restrict our analysis to individuals who lived in the same zip code in 2010 and 2012 and for whom moving to a different part of Los Angeles thus was not a

[7] For readability, we suppress the coefficients on all the control variables in table 6 and other tables in the main body of the paper. The appendix presents all coefficients on control variables from the main specifications in these tables.

TABLE 6
EFFECTS ON PROBABILITY OF HOME OWNERSHIP

	Dependent Variable: Pr(Owner in 2012)				
	(1)	(2)	(3)	(4)	(5)
	A. 2010 Renters				
Δ friend house prices, 2008–10 (%)	.608***	.511***	.501***	.544***	.672***
	(.042)	(.044)	(.169)	(.043)	(.043)
Δ friend county income, 2008–10 (%)				.332***	
				(.033)	
Δ friend house prices, 2010–12 (%)					.324***
					(.044)
Zip 2010 × zip 2012 fixed effects	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes
Sample restriction		Stayed in same zip code	Geographically nonclustered professions		
Observations	433,836	302,686	433,836	433,836	433,836
R^2	.43	.13	.43	.43	.43
Mean dependent variable	17.8	10.3	17.8	17.8	17.8
	B. 2010 Owners				
Δ friend house prices, 2008–10 (%)	.201***	.092***	.088***	.190***	.221***
	(.015)	(.013)	(.032)	(.016)	(.016)
Δ friend county income 2008–10 (%)				.049***	
				(.011)	
Δ friend house prices, 2010–12 (%)					.095***
					(.016)
Zip 2010 × zip 2012 fixed effects	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes
Sample restriction		Stayed in same zip code	Geographically nonclustered professions		
Observations	1,035,523	892,250	1,035,523	1,035,523	1,035,523
R^2	.56	.14	.56	.56	.56
Mean dependent variable	93.5	98.1	93.5	93.5	93.5

NOTE.—The table shows results from regression (5). The sample consists of Facebook users who lived in Los Angeles County in 2010 and whom we can match across the 2010 and 2012 Acxiom snapshots (the "change-of-tenure sample"). The dependent variable is an indicator capturing whether the individual is a homeowner in 2012. Panel A focuses on individuals who were renting in 2010 and panel B on individuals who were owning in 2010. All specifications control for 2010 and 2012 zip code–pair fixed effects, as well as demographic characteristics of the individuals; coefficients on the control variables are presented in the appendix. The house price experiences of all friends are instrumented for by the house price experiences of out-of-commuting-zone friends. Column 2 shows results for individuals who stayed in the same zip code. Column 3 exploits variation in friends' house price experiences only among individuals who are retired or work in geographically nonclustered professions. Column 4 adds the average income changes in friends' counties between 2008 and 2010 as a separate control. Column 5 adds friends' house price experiences between 2010 and 2012 as a separate control. Standard errors are clustered at the 2010 zip code level.

* Significant at $p < .10$.
** Significant at $p < .05$.
*** Significant at $p < .01$.

driver of ownership change. The average probability of transitioning from renting to owning is somewhat lower in this sample, at 10.3 percent. The estimated effect of friends' house price experiences on the probability of buying a home is only marginally smaller than in the full sample.

The specifications in columns 3 and 4 of table 6 address concerns that our results could be driven by common shocks to individuals and their friends that are large enough to move aggregate house prices in those geographically distant regions where these friends live. Column 3 exploits variation in friends' house price experiences only among individuals who either are retired or work in geographically nonclustered professions. The estimated effect of friends' house price experiences is only marginally smaller.[8] In column 4, we directly control for the recent income changes in the counties where the individual has friends. A 1 percentage point higher average income growth among a renter's friends is associated with a 0.33 percentage point higher probability of buying a house over the next 2 years. This could, for example, be picking up an effect of people working in professions that are disproportionately prevalent in counties where they have friends. Importantly for us, the estimated effect of friends' house price experiences on the purchasing decision is nearly unchanged.[9] These findings suggest that our results are not driven by unobserved income shocks.

So far, we have focused on the effects of friends' house price experiences between 2008 and 2010 on a 2010 renter's decision of whether to buy a house by 2012. However, friends' house price experiences between 2010 and 2012 might also affect the decision to buy a house by 2012. In column 5, we therefore present results from a regression that also controls for these experiences. A 1 percentage point higher house price experience of an individual's friends between 2010 and 2012 indeed further increases the likelihood of that individual becoming a homeowner by 2012 by about 0.32 percentage points. The effect of friends' house price experiences between 2008 and 2010 is unaffected.

We now turn to the extensive margin decisions of 2010 homeowners. Specifically, panel B of table 6 explores how friends' house price experiences affect the probability that 2010 homeowners sell their home by 2012. Only about 6 percent of 2010 homeowners become renters by 2012. The results suggest that homeowners whose friends experienced particu-

[8] In this specification, we add an indicator that is equal to one for all professions not identified as geographically nonclustered and set FriendHPExp equal to zero for these individuals. This allows us to exploit variation in FriendHPExp coming only from individuals in geographically nonclustered professions, while using the full sample to estimate the effect of the control variables and fixed effects.

[9] The correlation of average house price growth and average income growth between 2008 and 2010 across the social networks in our sample is 37 percent. This shows that there is substantial variation that allows us to separately identify the effect of house price changes and income changes in the geographies where an individual has friends.

larly large house price declines are more likely to sell their house. For 2010 homeowners, the magnitude of the effect of friends' house price experiences on the probability of owning a home in 2012 is a quarter to a third of the magnitude for 2010 renters.

B. *Social Networks and Housing Markets: Transaction Analysis*

Sample Description and Summary Statistics

In addition to analyzing the probability of individuals transitioning between renting and owning across the 2010 and 2012 Acxiom snapshots, we investigate how other dimensions of their housing investment decisions are affected by their friends' house price experiences. To do this, we use the fact that we observe information on all housing transactions since 1993 that led to an ownership spell that was ongoing as of either Acxiom snapshot. We can match more than 520,000 of such housing transactions in Los Angeles county to the social networks of the respective home buyers. We refer to this sample of transactions as the "transaction sample."

Table 7 provides summary statistics on the transactions and home buyers in the transaction sample; the appendix provides additional summary statistics. The average transaction price was \$403,344, and the average loan-to-value ratio at origination was about 85 percent.[10] The average property size was 1,775 square feet. The average home buyer was 35 years old at the time of the transaction and has 408 total friends and 156 out-of-commuting-zone friends.[11] As before, we observe substantial variation in friends' house price experiences across buyers who purchased properties at the same point in time: after conditioning on the transaction quarter, the across-buyers standard deviation of $\text{FriendHPExp}_{i,t-24m,t}$ is about 3.5 percent (see panels C and D of fig. 5 for the full distribution).

Property Size Results

We next analyze whether, conditional on buying a house, the house price experiences of a buyer's friends affect the intensive margin of her property investment. The unit of observation in regression (6) is a purchase of

[10] We observe transaction prices and mortgage amounts in ranges of about \$50,000. We take the midpoints of these ranges as the transaction price and mortgage amount.

[11] For some of the transactions, a property is purchased by more than one individual, and we can match both individuals to their Facebook accounts. In these cases, we average the set of demographic characteristics and pool the friends of the two buyers in our calculation of friends' house price experiences. Considering only the characteristics and friends' house price experiences of the head of household yields very similar results. Observing multiple buyers for the same transaction is the main reason why we have fewer observations in the transaction sample than we observe 2010 owners.

TABLE 7

SUMMARY STATISTICS: TRANSACTION SAMPLE

	Mean	Standard Deviation	Standard Deviation | Q	P5	P25	P50	P75	P95
Transaction characteristics:								
Transaction price (US$)	403,344	293,533	263,173	125,000	175,000	325,000	550,000	900,000
Origination LTV (%)	85.44	17.26	17.07	50.00	73.33	84.62	100.00	113.64
Property characteristics:								
Is single-family residence	.77	.42	.42	0	1	1	1	1
Property size (sq. ft.)	1,775	870	868	860	1,217	1,566	2,107	3,347
Lot size (sq. ft.)	9,452	9,374	9,302	2,500	7,500	7,500	7,500	25,000
Age of property (years)	40.5	24.7	24.5	1	21	43	56	83
Has pool	.23	.42	.42	0	0	0	0	1
Buyer characteristics:								
Number of friends:								
All friends	408	503	502	42	117	245	502	1,293
Out-of-commuting-zone friends	156	262	262	14	34	74	170	558
Δ friend house prices past 24 months (%):								
All friends	7.7	20.1	3.5	−31.1	−5.7	10.6	21.3	39.1
Out-of-commuting-zone friends	6.6	15.5	4.0	−24.3	−2.2	9.4	17.0	28.3
Age at purchase	35.4	14.3	14.1	18	28	35	43	58
Has high school diploma in 2010	.43	.50	.49	0	0	0	1	1
Has college degree in 2010	.38	.49	.49	0	0	0	1	1
Has graduate degree in 2010	.2	.4	.4	0	0	0	0	1
Income in 2010 ($1,000s)	79.5	41.2	40.8	17.5	45.0	62.5	112.5	150.0
Household size in 2010	3.1	1.7	1.6	1	2	3	4	6
Married in 2010	.6	.5	.5	0	0	1	1	1

NOTE.—The table shows summary statistics for the transaction sample, which consists of all housing transactions by Facebook users in Los Angeles County between 1993 and 2012 that led to a home ownership spell that was still ongoing as of the 2010 or 2012 Acxiom snapshots. $N = 526{,}594$. We present details on the transaction, the property that was transacted, and the buyer. For each characteristic, we show the mean, standard deviation, within-quarter standard deviation, and percentiles of the distribution.

property h by individual i at time t. The dependent variable is the log square footage of the property, multiplied by 100 to ease interpretation of the coefficients. The key explanatory variable, $\text{FriendHPExp}^{\text{All}}_{i,t-24m,t}$, is constructed as in equation (2) and captures the average house price changes experienced by buyer i's friends in the 24 months prior to the purchase:

$$\log(\text{PropSize}_{h,i,t}) = \alpha + \beta\text{FriendHPExp}^{\text{All}}_{i,t-24m,t} + \gamma\mathbf{X}_{i,2010} + \psi_t + \epsilon_{h,i,t}. \quad (6)$$

Table 8 presents estimates from regression (6). In column 1, we control for purchase-month fixed effects, ψ_t, and buyer characteristics, $\mathbf{X}_{i,2010}$.[12] The estimates suggest that a 5 percentage point (1.4 within-quarter standard deviations) increase in friends' average house price experiences is correlated with buyers purchasing a 1.6 percent larger property. This shows that individuals purchase larger properties when their friends have experienced more positive recent house price changes. To put the magnitude of the effect into perspective, a one standard deviation increase in the house price experiences of an individuals' friends has the same effect on the size of the purchased property as a $3,000 increase in annual household income.

Columns 2–6 address a number of potential concerns with our causal interpretation of the estimates of β in regression (6). For property purchases before 2010, we observe information on the transaction only if the property does not get resold prior to 2010. If the probability of a fast resale was correlated with both house price experiences of the buyers' friends and the size of the house bought, this selection could bias our results. In column 2, we therefore focus on sales since 2010, for which we observe a nonselected sample. The point estimate of β in this sample is slightly larger than the point estimate in the full sample, though the two estimates are not statistically distinguishable. This suggests that the selection of our transaction sample does not bias the results.

A second concern with our interpretation of β was that even though we exploit variation in the house price experiences only of friends living outside of Los Angeles, the capital gains or own experiences of buyers who recently moved from these regions to Los Angeles might still be correlated with the experiences of those friends. To test whether this confounds our estimates, column 3 restricts the sample to purchases since 2010 for which we can verify that the buyer lived in Los Angeles in 2010. The effects are nearly identical to those in the sample of all purchases since 2010. Housing wealth effects or an extrapolation of own house price experiences thus cannot explain our findings.

[12] We observe buyer age at the time of the transaction, but for other buyer characteristics, such as occupation, marital status, and household size, we use values from the most proximate Acxiom snapshot.

TABLE 8
Effects on Size of Property Purchased

	Dependent Variable: 100 × Log(Property Size)					
	(1)	(2)	(3)	(4)	(5)	(6)
Δ friend house prices past 24 months (%)	.310***	.520***	.530***	.400***	.234***	.285***
	(.053)	(.144)	(.164)	(.060)	(.079)	(.056)
Δ friend county income past 24 months (%)						.173***
						(.096)
Month fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Buyer controls	Yes	Yes	Yes	Yes, ×year	Yes	Yes
Sample notes		Purchases since 2010	Purchases since 2010, lived in LA in 2010		Geographically nonclustered professions	
Observations	526,594	95,561	68,388	526,593	526,594	526,594
R^2	.194	.134	.126	.204	.194	.194

Note.—The table shows results from regression (6). The sample consists of all housing transactions by Facebook users in Los Angeles County between 1993 and 2012 that led to a home ownership spell that was still ongoing as of the 2010 or 2012 Acxiom snapshots (the "transaction sample"). The dependent variable is the log of property size, multiplied by 100. The house price experiences of all friends are instrumented for by the house price experiences of out-of-commuting-zone friends. All columns control for purchase-month fixed effects and buyer characteristics; coefficients on the control variables are presented in the appendix. Column 2 restricts the sample to home purchases since 2010; col. 3 further restricts this to purchases since 2010 in which the buyers lived in Los Angeles County in 2010. In col. 4, we interact buyer characteristics with year-of-transaction fixed effects. Column 5 exploits variation in friends' house price appreciation only among buyers who are retired or work in geographically nonclustered professions. Column 6 adds the average income changes in the friends' counties in the 24 months prior to the transaction as a separate control. Standard errors are clustered at the purchase-month level.

* Significant at $p < .10$.

** Significant at $p < .05$.

*** Significant at $p < .01$.

A further challenge to our causal interpretation of β comes from characteristics of the buyers that might have a particularly strong direct effect on property investments in years when the buyers' geographically distant friends experience significant house price increases. In column 4, we limit the scope of such possible confounding effects by interacting buyer characteristics with purchase-year fixed effects. In column 5, we exploit variation in FriendHPExp only among buyers who are retired or work in geographically nonclustered professions. In column 6, we include direct controls for income changes in the buyers' social networks. Across these specifications, the estimates of β are similar to our baseline estimates, indicating that common shocks to individuals and their social networks do not explain the observed effect of friends' house price experiences.

Transaction Price Results

In this section, we analyze the effects of the house price experiences of both the buyers' friends and the sellers' friends on transaction prices. Conceptually, the property valuations of both buyers and sellers could be affected by their friends' house price experiences. In any bargaining model, the final transaction prices will then vary with these valuations.

For this analysis, we again consider the transaction sample and run hedonic regression (7). The unit of observation is the purchase of property h by individual i at time t. The dependent variable is the log of the transaction price, multiplied by 100 to ease the interpretation of the coefficients. All specifications include zip code × transaction year fixed effects, $\phi_{\text{zip}} \times \psi_{y(t)}$, allowing us to nonparametrically control for different time trends in prices across zip codes. We also control for buyer characteristics, $\mathbf{X}_{i,2010}$, and for property characteristics, $\mathbf{Z}_h$:

$$\log(\text{Price}_{h,i,t}) = \alpha + \beta \text{FriendHPExp}^{\text{All}}_{i,t-24m,t} + \delta \mathbf{X}_{i,2010} + \gamma \mathbf{Z}_h + \phi_{\text{zip}} \times \psi_{y(t)} + \epsilon_{h,i,t}. \tag{7}$$

Panel A of table 9 presents the main results from regression (7). The estimate in column 1 suggests that when home buyers' friends experience a 5 percentage point higher house price appreciation, the transaction price for a given home is 2.3 percent higher. To put this magnitude into perspective, it approximately corresponds to the price difference between a 1,140 square foot property and a 1,200 square foot property. The R^2 of the regression is over 80 percent, confirming that our hedonic property characteristics capture many of the important determinants of house prices.

While the hedonic regression controls for many determinants of property value, one might be concerned that individuals with larger house price increases in their social networks purchase properties that differ on unobservable characteristics, which could bias the estimates of β in regression (7). To rule out such confounding effects, column 2 includes property fixed effects in the regression. In this specification, β is identified only by transactions of properties for which we observe two transactions.[13] Since we are comparing transaction prices for the same property, this specification holds constant all unobservable characteristics of the properties. Overall, we observe 34,732 transactions for properties that trade twice in our sample. As one would expect, including property fixed effects increases the R^2 further, to 95 percent. Reassuringly, the effect of positive house price experiences among a buyer's friends on the transaction price is unaffected.

[13] In order to identify such repeat sales of the same property, one of the transactions has to occur before 2010 and the other between 2010 and 2012, so that we see the property attached to a different owner across the two Acxiom snapshots.

TABLE 9
EFFECTS ON TRANSACTION PRICE

	Dependent Variable: 100 × Log(Price)				
	(1)	(2)	(3)	(4)	(5)
			A. Main Results		
Δ friend house prices, buyer—past 24 months (%)	.452***	.486***	.408***	.445***	.335***
	(.015)	(.050)	(.076)	(.015)	(.068)
Δ friend house prices, seller—past 24 months (%)				.233***	.289**
				(.059)	(.112)
Year × zip code fixed effects, controls	Yes	Yes	Yes	Yes	Yes
Sample or specification notes		Property FE	Buyer FE		Property FE
Observations	523,299	34,732	32,226	523,299	33,230
R^2	.808	.950	.948	.809	.956
			B. Robustness Checks		
Δ friend house prices, buyer—past 24 months (%)	.468***	.461***	.489***	.423***	.441***
	(.029)	(.032)	(.016)	(.043)	(.015)
Δ friend county income, buyer—past 24 months (%)					.128***
					(.043)
Year × zip code fixed effects, controls	Yes	Yes	Yes	Yes	Yes
Sample or specification notes	Purchases since 2010	Purchases since 2010, lived in LA in 2010	Buyer controls × year	Geographically nonclustered professions	
Observations	95,226	68,202	523,299	523,299	523,299
R^2	.842	.847	.809	.810	.808

NOTE.— The table shows results from regression (7). The sample consists of all housing transactions by Facebook users in Los Angeles County between 1993 and 2012 that led to a home ownership spell that was still ongoing as of the 2010 or 2012 Acxiom snapshots (the "transaction sample"). The dependent variable is the log of the transaction price, multiplied by 100. The house price experiences of all friends are instrumented for by the house price experiences of out-of-commuting-zone friends. All columns control for year-of-purchase × zip code fixed effects, as well as property and buyer characteristics; coefficients on the control variables are presented in the appendix. In panel A, cols. 2 and 5 also include property fixed effects, and col. 3 includes buyer fixed effects. Columns 4 and 5 include the house price experiences of the sellers' social networks, instrumented for by their out-of-commuting-zone counterparts. Panel B presents robustness checks. Column 1 restricts the sample to house purchases since 2010; col. 2 further restricts this to purchases since 2010 in which the buyers lived in Los Angeles County in 2010. In col. 3, we interact buyer characteristics with year-of-transaction fixed effects. Column 4 exploits variation in friends' house price appreciation only among buyers who are retired or work in geographically nonclustered professions. Column 5 adds the average income changes in the buyers' friends' counties in the 24 months prior to the transaction as a separate control. Standard errors are clustered at the zip code level.

* Significant at $p < .10$.
** Significant at $p < .05$.
*** Significant at $p < .01$.

One concern with our causal interpretation of β is that friends' house price experiences might be correlated with unobserved buyer characteristics that could have a direct effect on her housing investment decisions. We argued that this was unlikely, since our identification comes from the interaction of the geographic dispersion of an individual's friends and the time-varying house price movements in those counties. To highlight this source of identification, column 3 includes buyer fixed effects. In this specification, all identification comes from individuals whom we observe purchasing more than one property.[14] Across those transactions, the friendship networks and unobservable characteristics of the buyers are held fixed, and the only force shifting FriendHPExp comes from the differential house price development in the fixed friendship networks prior to the two points in time when the individuals bought their homes. Our estimate of β is very similar in this specification, highlighting that our results are not confounded by the correlation between individuals' demographic characteristics and the geographic dispersion of their social networks.

So far, we have focused on the effect of the house price experiences of the buyers' friends on the transaction price. In columns 4 and 5 of table 9, we include the house price experiences of the sellers' friends in the 24 months before the sale as an additional regressor.[15] When sellers' friends experience a 5 percentage point higher house price appreciation, the transaction price is 1.2 percent higher. The estimated effect of the buyers' friends' house price experiences is similar to our baseline estimates. This evidence is consistent with friends' house price experiences affecting both the buyers' and the sellers' valuations of the property and, thus, their reservation prices in the bargaining that determines the transaction price.

The specifications in panel B of table 9 correspond to those in columns 2–6 of table 8 and address a number of challenges to our identifi-

[14] In order to identify two transactions by the same buyer, we need the same individual to owner-occupy two different properties in the 2010 and 2012 Acxiom snapshots. This will then allow us to observe information about the transaction that initiated each ownership spell.

[15] We observe information on the seller only for transactions between 2010 and 2012, since for those transactions we know who owned the property in 2010, prior to its being sold. We can match the sellers in about 20,000 transactions to their Facebook accounts. We include all transactions in the regression, even if we cannot match the seller to Facebook, in order to increase the power for estimating the coefficients on the property characteristics and the buyer experience. In particular, in that specification, we also include an indicator, FBMiss_i, that is equal to one for all transactions in which we cannot match the seller to his or her Facebook profile, and zero otherwise. We set FriendHPExp of the seller equal to zero when FBMiss_i equals one. However, estimates are similar if we focus only on the transactions for which we can identify both the buyer and the seller.

cation of causal effects. As was the case in our analysis of the intensive margin decision, we do not find evidence for a number of noncausal alternative explanations of the patterns in the data.

C. *Further Robustness Checks and Differential Effects*

In the appendix, we provide additional robustness checks to our results. We also explore a number of sample splits to analyze whether effects differ across the population. First, we show that results are similar when we use the house price experiences of out-of-state friends as an instrument instead of the house price experiences of out-of-commuting-zone friends. Second, we use friends' house price experiences over the prior 12 months, 36 months, and 48 months as explanatory variables instead of the house price experiences over the prior 24 months, as we do in the main body of the paper. The magnitude of the effect is declining as we increase the time window over which we measure friends' house price experiences, suggesting that the most recent experiences within a person's social network have the largest effects on her behavior. Third, we analyze the effects separately for individuals in different age groups and with different education levels. The effects of friends' house price experiences on the probabilities of buying a house for renters or selling a house for owners are declining in the age of the individuals. The effects of these experiences on the size of the property purchased, and the price paid for a given property, are stable across age groups. There are no systematic differences in the effect sizes across individuals with different education levels.

We also consider whether the response of individuals' housing investment behavior to their friends' house price experiences is different during periods with booming housing markets relative to periods with more stable or declining housing markets. The effect on the price paid is nearly identical across the housing boom period 2001–6, the housing bust period 2007–9, and the relatively flat period 2010–12. The effect on size bought is somewhat larger in the boom and flat periods than it is in the housing bust period. These findings suggest that the social dynamics channel we document in this paper is likely to be active during both housing booms and busts.

IV. Social Networks and Housing Markets: Mechanisms

In the previous section, we documented a causal relationship between the house price experiences in an individual's social network and her housing investment behavior. In this section, we investigate potential mechanisms behind this causal relationship.

A. *Evidence for Expectations Channel*

A first plausible channel through which the house price experiences in a person's social network can affect her housing market investments is by influencing her perceptions of the attractiveness of local property investments. A number of possible mechanisms for such an expectations-based channel have been proposed in the literature, though little systematic empirical evidence has been developed.

One important dimension along which existing theories differ is whether agents are mechanically influenced by signals they receive through their social networks, or whether they attempt to rationally extract information from the experiences of their friends. For example, one prominent narrative of the role of social interactions in the housing market has been put forward by Shiller (2008b, 96), who writes that "many people seem to be accepting that the recent home price experience is at least in part the result of a social epidemic of optimism for real estate." Shiller (2008a) describes "the contagious optimism, seeming impervious to facts, that often takes hold when prices are rising. . . . Speculative bubbles are fueled by the social contagion of boom thinking." In Shiller's narrative, individuals are mechanically "infected" by the beliefs of their friends, which, in turn, are driven by these friends extrapolating from their own house price experiences. This response to friends' beliefs is independent of whether friends' house price experiences contain useful information for predicting own local house price changes.

We next document that social interactions indeed have important effects on individuals' perceptions of the attractiveness of local housing market investments. We also attempt, to the extent possible, to understand whether our findings are more consistent with a mechanical or a rational change in beliefs in response to friends' house price experiences.

Description of Expectations Survey

Investigating the response of an individual's expectations to her friends' house price experiences presents an additional challenge: how do we measure those expectations? To overcome this measurement challenge, we analyze responses to a short user survey conducted by Facebook in November 2015. The survey targeted Facebook users living in a few Los Angeles zip codes through a post on their News Feed.[16] It informed users that "Facebook is helping researchers understand what real people think about the economy. Your survey responses will be combined with

[16] A person's News Feed is a personalized, constantly updating list of content posted by friends and followed pages (e.g., messages, photos, videos), advertisements, and surveys. It is shown to users as the landing page when they log on to Facebook. The appendix shows the survey interface.

the information that you publicly share on Facebook and average house prices to better help us understand the housing economy. Help us out by answering the following questions, your responses will be kept anonymous," followed by four multiple-choice questions:

1. How informed are you about house prices in your zip code?
   □ Not at all informed □ Somewhat informed □ Well informed
   □ Very well informed
2. How informed are you about house prices where your friends live?
   □ Not at all informed □ Somewhat informed □ Well informed
   □ Very well informed
3. How often do you talk to your friends about whether buying a house is a good investment?
   □ Never □ Rarely □ Sometimes □ Often
4. If someone had a large sum of money that they wanted to invest, would you say that relative to other possible financial investments, buying property in your zip code today is:[17]
   □ A very good investment □ A somewhat good investment
   □ Neither good nor bad as an investment
   □ A somewhat bad investment □ A very bad investment

We observe 1,242 survey responses: 55 percent of respondents are male, and their ages range between 19 and 75 years, with an average of 46 years and an interquartile range of 35–56 years. Respondents come from 113 Los Angeles zip codes, but 40 percent of them live in the 20 most-represented zip codes.[18]

Panels A–D of figure 6 plot the distribution of responses to each survey question. Most respondents believe that buying property is at least a somewhat good investment, but we observe significant heterogeneity in respondents' beliefs about the attractiveness of local real estate investments. About 73 percent of individuals claim to be at least "somewhat informed" about house prices where their friends live, while 27 percent are "well informed" or "very well informed." Over half of the respondents report talking at least "sometimes" to their friends about whether buying property is a good investment, while 15 percent talk "often." This suggests a potentially important role for social interactions in influencing housing market expectations and investments. There is no relation-

[17] The wording to this question corresponds to a question on the New York Fed Survey of Consumer Expectations.

[18] As is generally the case with analyzing survey data, there is some concern that individuals who respond to a survey might be different on important characteristics. While we found that respondents look similar to the targeted population on observable characteristics, it could be, e.g., that those individuals who respond to a housing expectation survey are disproportionately likely to talk to their friends about housing investments.

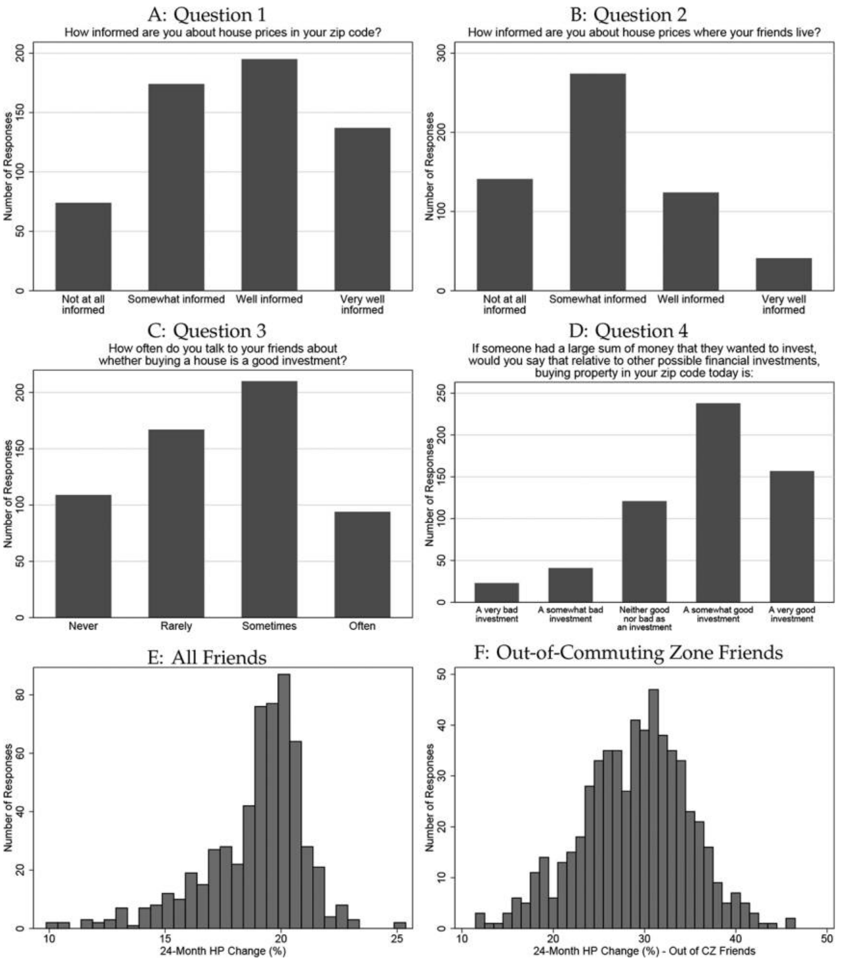

Fig. 6.—Expectations survey. Panels A–D present the distribution of the 1,242 responses to the expectations survey conducted by Facebook in November 2015. We analyze and describe this survey in Section IV.A. Panels E and F provide the average house price experience in the survey respondents' social network in the 24 months prior to taking the survey, both for all friends and for all friends living outside the Los Angeles commuting zone.

ship between an individual's friends' house price experiences and her propensity to talk to friends about investing in the housing market: the average house price experiences of the respondents' friends, split up by their responses to question 3, are 18.4 percent, 18.3 percent, 18.3 percent, and 18.5 percent, respectively.

We next analyze how the average house price movements in individual i's social network in the 24 months before answering the survey, $\text{FriendHPExp}^{\text{All}}_{i,2013,2015}$, affect her optimism about property investments in her own zip code. There is significant across-respondent variation in this experience measure, which has a mean of 18.3 percent, a standard deviation of 2 percent, and a 10–90 percentile range of 4.5 percent. Panels E and F of figure 6 plot the distributions across the survey respondents of the house price experiences of their friends and their out-of-commuting-zone friends.

Analysis of Expectations Survey

Regression (8) analyzes the relationship between the house price experiences of individuals' friends and their beliefs about whether buying local property is a good investment. The dependent variable is the individuals' response to question 4.[19] The vector $\mathbf{X}_i$ controls for the age and gender of the respondent. Since respondents are asked to evaluate the attractiveness of buying property in their own zip code and the true attractiveness of such investments can vary across zip codes, we also include zip code fixed effects, ψ_{zip}:

$$\text{ResponseQ4}_i = \alpha + \beta\text{FriendHPExp}^{\text{All}}_{i,2013,2015} + \gamma\mathbf{X}_i + \psi_{\text{zip}} + \epsilon_i. \quad (8)$$

To deal with the ordinal nature of the responses to question 4, we code the answers as 1–5, with 5 corresponding to the most optimistic view on property investments. This approach assumes that the "distance" between each of the five possible answers to question 4 is the same.[20] The resulting measure of optimism about property investments has a standard deviation of 1.06. Most of this heterogeneity is across individuals responding about investing in property in the same zip code: when conditioning on ψ_{zip}, the residual standard deviation of ResponseQ4_i remains at 0.98. As before, we estimate regression (8) using an IV strategy, where we instrument for the house price experiences of all friends with the house price experiences of only the out-of-commuting-zone friends.

Column 1 of table 10 presents estimates of equation (8). Holding zip code, age, and gender fixed, an increase in friends' house price experiences makes respondents more optimistic about the attractiveness of investing in property. Quantitatively, a one standard deviation increase in $\text{FriendHPExp}^{\text{All}}_{i,2013,2015}$ is associated with a statistically significant 0.08 standard deviation increase in our measure of optimism. It is difficult to assess

[19] Our favorite interpretation of the responses to this question is that they reflect differences in the physical probabilities that respondents assign to different states of the world. However, it is possible that respondents risk-adjust their answers to whether they think that housing is an attractive investment. In that case, different responses could also reflect differences in risk adjustments of respondents whose friends experience different house price movements.

[20] In the appendix, we also take a second approach to dealing with the ordinal nature of the responses to question 4. In particular, we present cumulative odds ratios from an ordered logit model. The conclusions from this specification are very similar to those in table 10.

TABLE 10

Effects on Belief of Whether Property Is a Good Investment

	Dependent Variable: Local Housing a Good Investment? (Question 4)				
	(1)	(2)	(3)	(4)	(5)
Δ friend house prices, 2013–15 (%)	.040**	.036*			
	(.017)	(.019)			
Δ friend house prices, 2013–15 (%) × ordering of question:					
Expectation question last			.039**		
			(.021)		
Expectation question first			.048*		
			(.029)		
Δ friend house prices, 2013–15 (%) × knowledge of house prices where friends live:					
Not at all informed				.002	
				(.036)	
Somewhat informed				.036	
				(.023)	
Well informed				.068*	
				(.039)	
Very well informed				.119*	
				(.069)	
Δ friend house prices, 2013–15 (%) × talk with friends about housing investments:					
Never					−.050
					(.038)
Rarely					.001
					(.028)
Sometimes					.086***
					(.027)
Often					.096**
Demographic controls	Yes	Yes	Yes	Yes	Yes
Zip code fixed effects	Yes	Yes	Yes	Yes	Yes
Sample		LA in 2012			
Observations	1,242	1,110	1,242	1,242	1,242

Note.—The table shows results from regression (8). The dependent variable is the answer to survey question 4: "If someone had a large sum of money that they wanted to invest, would you say that relative to other possible financial investments, buying property in your zip code today is: (1) A very bad investment, (2) A somewhat bad investment, (3) Neither good nor bad as an investment, (4) A somewhat good investment, or (5) A very good investment," with the five (ordered) answers recoded to 1–5. The house price experiences of all friends are instrumented for by their out-of-commuting-zone counterparts. Column 1 shows the baseline estimates. Column 2 restricts the sample to respondents who lived in Los Angeles in 2012. The last three columns estimate differential effects by the ordering of the questions (col. 3), by how informed respondents claimed to be about house prices where their friends live (col. 4), and by how often they reported talking to their friends about investing in property (col. 5). The specifications in cols. 3, 4, and 5 also include non-interacted indicator variables for the question ordering and the possible responses to questions 2 and 3, respectively; in the interest of space, the corresponding coefficients are not reported. All columns also control for respondent age and gender. Standard errors are in parentheses.

* Significant at $p < .10$.

** Significant at $p < .05$.

*** Significant at $p < .01$.

the magnitude of this effect. However, in recent work, Armona, Fuster, and Zafar (2016) have analyzed whether an individual's perceived past house price changes in her own zip code affect her beliefs about the attractiveness of housing investments, elicited through a survey question identical to question 4 above. They find that a one standard deviation increase in individuals' perceptions of their own local house price changes over the past year is associated with an increase in ResponseQ4_i of about 0.1. They compare these estimates to our findings and conclude that the effect on expectations through the social dynamics channel that we highlight in this paper is of a magnitude similar to that of the effect through the extrapolation from past local house price movements that is the focus of their research.

For survey respondents who only recently moved to Los Angeles, possibly from areas in which they have many friends, $\text{FriendHPExp}^{\text{OutCZ}}_{i,2013,2015}$ might be correlated with their own house price experience, in which case even the IV strategy could not separate the effect of social interactions from that of extrapolative expectations. In column 2, we thus restrict the sample to survey respondents who already lived in Los Angeles in 2012. The results in this subsample are very similar.

A common concern with analyzing survey data is the possibility that the framing and ordering of questions affect the responses. In particular, given the order of questions described above, one might worry that by first asking respondents whether they knew about house prices where their friends live, one might prime them to place more weight on those friends' experiences when subsequently reporting their own perceptions of the attractiveness of housing market investments. To rule out such effects, for about 35 percent of respondents the order of questions was reversed, asking them first about their housing market expectations before eliciting responses to the other questions. Column 3 shows that the correlation between a respondent's friends' house price experiences and her own expectations is, if anything, slightly stronger in the sample of respondents who first reported their housing market beliefs. This suggests that framing effects do not significantly affect our results.

We next provide additional evidence that the correlation between an individual's housing market beliefs and her friends' house price experiences is driven by social interactions and not by other confounding shocks. In column 4, we interact $\text{FriendHPExp}^{\text{All}}_{i,2013,2015}$ with each possible response to question 2. We also include noninteracted indicator variables for each possible response to question 2 but, in the interest of space, do not report the corresponding coefficients. The relationship between an individual's assessment of whether buying property is a good investment and the house price experiences of her friends is stronger for individuals who report being aware of house prices where their friends live. Similarly, in column 5 we interact $\text{FriendHPExp}^{\text{All}}_{i,2013,2015}$ with each possible response to question 3. For respondents who report that they regularly talk to their friends about whether buying property is a good investment, we find a strong relationship between their friends' house price experiences and their own assessment of whether property in their own zip code is a good investment. Indeed, for respondents who sometimes or often talk to their friends about property investments, the effect size is twice as large as the effect size for the average individual. For respondents who rarely or never talk to their friends about investing in the housing market, no statistically significant relationship is found. This finding suggests that the observed correlation is driven by social interactions, and not, for example, by people reading local newspapers from areas where they have friends.

Overall, these results suggest an important role for social interactions in affecting individuals' assessments of the attractiveness of housing market investments. All else equal, an individual perceives property to be a more attractive investment when there are larger house price gains within her social network. These effects are statistically significant, economically large, and more pronounced for those respondents who report talking with their friends about whether housing is a good investment.

B. *Reasons for Updating Expectations*

Why would an individual's perceptions of the attractiveness of local housing market investments be influenced by the house price movements in those geographically distant areas where she has friends? We next present evidence that can help differentiate between various possible explanations. In particular, we analyze whether we can find evidence in favor of a rational learning story.

A first such story would involve individuals learning through their friends about house price changes elsewhere that are predictive of future Los Angeles house price changes. To test whether such an effect contributes to our findings, panel A of table 11 splits individuals into groups based on how predictive the house price experiences of their out-of-commuting-zone social networks are for subsequent Los Angeles house price changes.[21] We then obtain separate estimates of the effect of friends' house price experience on the housing investment behaviors of each group. There is no evidence that people with more predictive social networks respond in a systematically different way.

[21] For every individual, we find the correlation between the house price movements in their out-of-commuting-zone social network over the previous 24 months and Los Angeles house price movements over the next 12 months. We estimate this correlation using yearly observations between 1993 and 2012. Varying the time horizons and the sample period does not significantly affect the ordering of individuals by the predictiveness of their social networks' house price experiences.

TABLE 11
DIFFERENTIAL EFFECTS BY NETWORK PREDICTIVENESS AND NETWORK SIZE

	Pr(Own in 2012) (1)	Pr(Own in 2012) (2)	100 × Log(Size) (3)	100 × Log (Price) (4)
	A. Effects by Predictiveness of Network			
Δ friend house prices past 24 months (%):				
Correlation < .5	.464***	.152***	.293***	.486***
	(.044)	(.017)	(.053)	(.017)
.5 ≤ correlation < .6	.552***	.185***	.313***	.457***
	(.043)	(.016)	(.050)	(.015)
Correlation ≥ .6	.461***	.176***	.299***	.445***
	(.049)	(.018)	(.049)	(.015)
Controls as in	Table 6, col. 1, 2010 renters	Table 6, col. 1, 2010 owners	Table 8, col. 1	Table 9, col. 1
p-value (high correlation == low correlation)	.946	.097	.660	.000
R^2	.434	.564	.204	.814
Observations	433,836	1,035,523	526,594	523,299
	B. Effects by Number of Counties in Friendship Network			
Δ friend house prices past 24 months (%):				
Below median number of counties	.605***	.159***	.316***	.450***
	(.057)	(.019)	(.054)	(.015)
Above median number of counties	.616***	.292***	.338***	.464***
	(.060)	(.028)	(.057)	(.016)
Controls as in	Table 6, col. 1, 2010 renters	Table 6, col. 1, 2010 owners	Table 8, col. 1	Table 9, col. 1
p-value (many counties == few counties)	.892	.000	.116	.103
R^2	.434	.564	.193	.808
Observations	433,836	1,035,523	526,594	523,299

NOTE.—The table shows results from the main instrumental variables regressions in tables 6, 8, and 9. In panel A, we analyze the effect of friends' house price experiences separately by the correlation between the house price movements in an individual's social network and subsequent Los Angeles house price movements (see n. 21 for details). In panel B, we analyze the effect of friends' house price experiences separately by the number of counties in which individuals have friends. Specifications, samples, and standard errors are as described in the original tables.
* Significant at $p < .10$.
** Significant at $p < .05$.
*** Significant at $p < .01$.

Similarly, many plausible rational explanations of the behavior we document involve individuals learning about some fundamental national housing demand shock from observing house price growth across multiple geographies. If this were an important channel, we would expect individuals to respond more to their friends' experiences if these friends were more geographically dispersed, since the average experience of these friends would then be more informative about the national shock. Yet panel B of table 11 shows that the response of individuals to the experiences of their friends does not generally vary with the number of counties these individuals have friends in.

These pieces of evidence point away from a rational learning explanation for our findings. This is perhaps unsurprising. Indeed, if the house price movements in a different part of the country were sufficiently informative to affect a rational agent's valuation of a given house by thousands of dollars, then, in a world of rational learning, everybody should update their expectation equally on the basis of these house price movements, which are available for free and in real time. We thus conclude that the evidence is most consistent with mechanical belief updating along the lines of Shiller (2008a, 2008b).

C. *Evidence against Alternative Causal Channels*

We next present evidence against three causal mechanisms other than social interactions through which friends' house price experiences could affect an individual's housing market behavior. Result tables and further discussions are provided in the appendix.

Bequests

A first alternative explanation is that the house price experiences in a person's social network may have a direct wealth or liquidity effect. In particular, if a person has many friends where her parents live, increases in house prices in that area might affect the value of any property owned by her parents. In that case, if this individual is expecting to inherit a more expensive house or if her parents have more resources to help her with purchasing a property in Los Angeles, this could influence her purchasing behavior through a channel that is unrelated to social dynamics.

We present three pieces of evidence that suggest this mechanism cannot explain our findings. First, we separately exploit variation in the overall social network house price experience coming from three subsets of out-of-commuting-zone friends: family members, work colleagues, and college friends. While an individual might expect higher future bequests when her family members experience higher house price growth, this is less likely to be the case for her college or work friends. Yet we find that the influence of the house price experiences in all three subnetworks on investment behavior is very similar, suggesting the bequest channel is relatively unimportant.

As a second piece of evidence against a bequest story, we show that our estimates are similar among individuals whose bequests are less likely to be affected by the house price movements of their US-based out-of-commuting-zone friends. In particular, the effects of friends' house price experiences on housing investments are quite similar when restricting the sample to individuals whose hometown is Los Angeles or to individuals whose hometown is outside of the United States.

Third, since most individuals can expect bequests from only a few close relatives within their social network, a bequest channel should be stronger when individuals' social networks are more geographically concentrated and when their friends' overall house price experience thus more closely corresponds to that of those close relatives. Yet we have shown that the effects are unrelated to how many counties an individual has friends in, providing further evidence against a bequest story.

Consumption Externalities

A second alternative explanation for our findings is the possible presence of consumption externalities across individuals and their friends. For example, an individual might buy a house to "keep up with the Joneses" after her friends purchased a home. Even though the construction of our key explanatory variable in equation (2) does not depend on whether an individual's friends have purchased a house, this does not completely alleviate the potential of consumption externalities to explain some of our findings. Indeed, since house prices and transaction volumes generally comove, people are more likely to buy a house, on average, in regions where house prices go up. Therefore, FriendHPExp could still be correlated with friends' home purchase behavior. However, when we directly include controls for the change and level of trading volume in the counties where an individual has friends, the estimated effects of friends' house price experiences are nearly identical, suggesting that they are not just picking up a desire to keep up with friends.

Hedging

A further alternative explanation of our results is that individuals plan to eventually move to those parts of the country where they have friends. In that case, when they see house prices there go up, they might want to purchase a house in Los Angeles in order to hedge against further (national) price increases that might price them out of the market. If this were an important force explaining our results, one would expect the effect to be larger for people whose friends live in housing markets that are more correlated with Los Angeles and for which Los Angeles housing would thus provide a better hedge. Yet we have shown that this is not the case.

V. Conclusion

In this paper, we highlight how newly emerging data from online social networking services allow researchers to better understand the economic effects of social interactions. To illustrate this point, we use anonymized administrative data from Facebook to document that the house price experiences within an individual's social network affect her perceptions of the attractiveness of property investments and through this channel have large effects on her housing market activity. Our results highlight that social interactions play an important role in determining how people form expectations as well as in explaining their actual investment behavior. The effects are quantitatively large and have the potential to affect aggregate outcomes. This suggests that, at the county level, friendship networks provide a mechanism that can propagate house price shocks through the economy. In related work, Bailey et al. (2018) show that other measures of economic activity, such as trade and patent citations, are also correlated with more aggregated social networks as measured by the Facebook social graph.

While we document the effect of social interactions on expectations and investment behavior in the housing market, it is likely that similar social dynamics are also at work in other settings. For example, it is possible that optimism and pessimism about stock market investments, or sentiments about the economy more generally, also spread through social interactions in a similar way. One interesting question left unexplored in this paper is whether the increasing connectedness of individuals through social media will itself have an effect on how the experiences of individuals influence the behavior of their friends across a number of settings. We hope that the increasing availability of data from online social networking services, such as the publicly available Social Connectedness Index described in Bailey et al. (2018), will facilitate more research along a number of these important dimensions.

References

Acemoglu, Daron, Munther A. Dahleh, Ilan Lobel, and Asuman Ozdaglar. 2011. "Bayesian Learning in Social Networks." *Rev. Econ. Studies* 78 (4): 1201–36.

Akerlof, George A., and Robert J. Shiller. 2010. *Animal Spirits: How Human Psychology Drives the Economy, and Why It Matters for Global Capitalism.* Princeton, NJ: Princeton Univ. Press.

Alatas, Vivi, Abhijit Banerjee, Arun G. Chandrasekhar, Rema Hanna, and Benjamin A. Olken. 2016. "Network Structure and the Aggregation of Information: Theory and Evidence from Indonesia." *A.E.R.* 106 (7): 1663–1704.

Angeletos, George-Marios, and Jennifer La'O. 2013. "Sentiments." *Econometrica* 81 (2): 739–79.

Armona, Luis, Andreas Fuster, and Basit Zafar. 2016. "Home Price Expectations and Behavior: Evidence from a Randomized Information Experiment." Staff Report no. 798, Fed. Reserve Bank New York.

Bailey, Michael, Ruiqing Cao, Theresa Kuchler, Johannes Stroebel, and Arlene Wong. 2018. "Social Connectedness: Measurement, Determinants, and Effects." *J. Econ. Perspectives* 32 (3): 259–80.

Bailey, Michael, Eduardo Davila, Theresa Kuchler, and Johannes Stroebel. 2017. "House Price Beliefs and Mortgage Leverage Choice." Working Paper no. 24091, NBER, Cambridge, MA.

Baker, Scott. 2018. "Debt and the Response to Household Income Shocks: Validation and Application of Linked Financial Account Data." *J.P.E.* 126 (4): 1504–57.

Barabási, Albert-László. 2016. *Network Science*. Cambridge: Cambridge Univ. Press.

Barberis, Nicholas, and Richard Thaler. 2003. "A Survey of Behavioral Finance." In *Handbook of the Economics of Finance*, edited by George M. Constantinides, Milton Harris, and Rene M. Stulz, 1053–1128. Amsterdam: Elsevier.

Bayer, Patrick, Christopher Geissler, Kyle Mangum, and James W. Roberts. 2016. "Speculative Fever: Investor Contagion in the Housing Bubble." Working Paper no. 22065, NBER, Cambridge, MA.

Berger, David, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra. 2017. "House Prices and Consumer Spending." Working Paper no. 21667, NBER, Cambridge, MA.

Breza, Emily, Arun G. Chandrasekhar, Tyler H. McCormick, and Mengjie Pan. 2017. "Using Aggregated Relational Data to Feasibly Identify Network Structure without Network Data." arXiv:1703.04157v2.

Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo. 2016. "Understanding Booms and Busts in Housing Markets." *J.P.E.* 124 (4): 1088–1147.

Choi, James J., David Laibson, Brigitte C. Madrian, and Andrew Metrick. 2009. "Reinforcement Learning and Savings Behavior." *J. Finance* 64 (6): 2515–34.

DeMarzo, Peter M., Dimitri Vayanos, and Jeffrey Zwiebel. 2003. "Persuasion Bias, Social Influence, and Unidimensional Opinions." *Q.J.E.* 118 (3): 909–68.

Duggan, Maeve, Nicole B. Ellison, Cliff Lampe, Amanda Lenhart, and Mary Madden. 2015. "Social Media Update 2014." Pew Research Center, Washington, DC.

Duggan, Maeve, and Dana Page. 2016. "Social Media Update 2016." Pew Research Center, Washington, DC.

Einav, Liran, Theresa Kuchler, Jonathan Levin, and Neel Sundaresan. 2015. "Assessing Sale Strategies in Online Markets Using Matched Listings." *American Econ. J.: Microeconomics* 7 (2): 215–47.

Enke, Benjamin, and Florian Zimmermann. 2017. "Correlation Neglect in Belief Formation." *Rev. Econ. Studies*. DOI:10.1093/restud/rdx081.

Facebook. 2017. "Facebook Form 10-Q, Quarterly Report, Q1." http://d1lge852tjjqow.cloudfront.net/CIK-0001326801/23221e2b-293c-453a-ac0f-50ad44ddf178.pdf.

Feld, Scott L. 1991. "Why Your Friends Have More Friends than You Do." *American J. Sociology* 96 (6): 1464–77.

Geanakoplos, John. 2009. "The Leverage Cycle." In *NBER Macroeconomics Annual*, vol. 24, edited by D. Acemoglu, K. Rogoff, and M. Woodford, 1–66. Chicago: Univ. Chicago Press.

Greenwood, Robin, and Andrei Shleifer. 2014. "Expectations of Returns and Expected Returns." *Rev. Financial Studies* 27 (3): 714–46.

Hampton, Keith, Lauren Sessions Goulet, Lee Rainie, and Kristen Purcell. 2011. "Social Networking Sites and Our Lives." Pew Internet and American Life Project, Washington, DC.

Harrison, J. Michael, and David M. Kreps. 1978. "Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations." *Q.J.E.* 92 (2): 323–36.

Hong, Harrison, and Jeremy C. Stein. 1999. "A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets." *J. Finance* 54 (6): 2143–84.

———. 2007. "Disagreement and the Stock Market." *J. Econ. Perspectives* 21 (2): 109–28.

Jackson, Matthew O. 2010. *Social and Economic Networks*. Princeton, NJ: Princeton Univ. Press.

Jackson, Matthew O., Brian Rogers, and Yves Zenou. 2017. "The Economic Consequences of Social Network Structure." *J. Econ. Literature* 55:49–95.

Jeffers, Jessica S. 2017. "The Impact of Restricting Labor Mobility on Corporate Investment and Entrepreneurship." Working paper. https://www.marshall.usc.edu/sites/default/files/2017-03/RECR_2-10-17_F_JEFFERS-the-impact-of-restricting-labor-mobility-on-corporate-investment-and-entrepreneurship.pdf.

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.

Kaplan, Greg, Kurt Mitman, and Giovanni L. Violante. 2017. "The Housing Boom and Bust: Model Meets Evidence." Working Paper no. 23694, NBER, Cambridge, MA.

Kogovšek, Tina, and Anuška Ferligoj. 2005. "Effects on Reliability and Validity of Egocentered Network Measurements." *Soc. Networks* 27 (3): 205–29.

Kuchler, Theresa. 2013. "Sticking to Your Plan: Hyperbolic Discounting and Credit Card Debt Pay-Down." Working Paper no. 24881, NBER, Cambridge, MA.

Kuchler, Theresa, and Basit Zafar. 2015. "Personal Experiences and Expectations about Aggregate Outcomes." Staff Report no. 748, Fed. Reserve Bank New York.

Landvoigt, Tim. 2017. "Housing Demand during the Boom: The Role of Expectations and Credit Constraints." *Rev. Financial Studies* 30 (6): 1865–1902.

Magno, Gabriel, Giovanni Comarela, Diego Saez-Trumper, Meeyoung Cha, and Virgilio Almeida. 2012. "New Kid on the Block: Exploring the Google+ Social Graph." In *Proceedings of the 2012 ACM Conference on Internet Measurement*, 159–70. New York: Assoc. Computing Machinery.

Malmendier, Ulrike, and Stefan Nagel. 2011. "Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?" *Q.J.E.* 126 (1): 373–416.

McPherson, Miller, Lynn Smith-Lovin, and James M. Cook. 2001. "Birds of a Feather: Homophily in Social Networks." *Ann. Rev. Sociology* 27:415–44.

Miller, Edward M. 1977. "Risk, Uncertainty, and Divergence of Opinion." *J. Finance* 32 (4): 1151–68.

Morris, Martina. 2004. *Network Epidemiology: A Handbook for Survey Design and Data Collection*. Oxford: Oxford Univ. Press on Demand.

Nathanson, Charles G., and Eric Zwick. 2017. "Arrested Development: Theory and Evidence of Supply-Side Speculation in the Housing Market." Working Paper no. 23030, NBER, Cambridge, MA.

Piazzesi, Monika, and Martin Schneider. 2009. "Momentum Traders in the Housing Market: Survey Evidence and a Search Model." *A.E.R.* 99 (2): 406–11.

Piazzesi, Monika, Martin Schneider, and Johannes Stroebel. 2017. "Segmented Housing Search." Working Paper no. 20823, NBER, Cambridge, MA.

Scheinkman, Jose A., and Wei Xiong. 2003. "Overconfidence and Speculative Bubbles." *J.P.E.* 111 (6): 1183–1220.

Shiller, Robert J. 2008a. "Infectious Exuberance." *Atlantic*, July/August.

———. 2008b. "Understanding Recent Trends in House Prices and Home Ownership." In *Housing, Housing Finance, and Monetary Policy: Proceedings of the Fed-*

eral Reserve Bank of Kansas City Economic Policy Symposium, 89–123. Kansas City: Fed. Reserve Bank Kansas City.

———. 2015. *Irrational Exuberance*. Princeton, NJ: Princeton Univ. Press.

Shin, Won-Yong, Bikash C. Singh, Jaehee Cho, and André M. Everett. 2015. "A New Understanding of Friendships in Space: Complex Networks Meet Twitter." *J. Information Sci*. 41 (6): 751–64.

Simsek, Alp. 2013. "Speculation and Risk Sharing with New Financial Assets." *Q.J.E*. 128 (3): 1365–96.

Ugander, Johan, Brian Karrer, Lars Backstrom, and Cameron Marlow. 2011. "The Anatomy of the Facebook Social Graph." arXiv:1111.4503.

Vissing-Jorgensen, Annette. 2003. "Perspectives on Behavioral Finance: Does 'Irrationality' Disappear with Wealth? Evidence from Expectations and Actions." In *NBER Macroeconomics Annual*, vol. 18, edited by Mark Gertler and Kenneth Rogoff, 139–208. Cambridge, MA: MIT Press.