Social Networks Shape Beliefs and Behavior: Evidence from Social Distancing during the COVID-19 Pandemic

Michael Bailey

Meta Platforms

Drew Johnston

Harvard University

Martin Koenen

Harvard University

Theresa Kuchler

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

Dominic Russel

Harvard University

Johannes Stroebel

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

Previously circulated as "Social Distancing during a Pandemic: The Role of Friends." This research was facilitated through a research consulting agreement between the academic authors and Meta. Bailey is an employee at Meta. We are thankful for helpful comments from many seminar and conference participants. This paper was edited by Lint Barrage.

Electronically published June 27, 2024

Journal of Political Economy Microeconomics, volume 2, number 3, August 2024.

https://doi.org/10.1086/729533

We analyze de-identified data from Facebook to show how social connections affect beliefs and behaviors in high-stakes settings. During the COVID-19 pandemic, individuals with friends in regions facing severe disease outbreaks reduced their mobility more than their demographically similar neighbors with friends in less affected areas. To explore why social connections shape behaviors, we show that individuals with higher friend exposure to COVID-19 are more supportive of social distancing measures and less likely to advocate to reopen the economy. We conclude that friends influence individuals' behaviors in part through their beliefs, even when there is abundant information from expert sources.

In the United States and many other countries, there is substantial public disagreement about important elements of established scientific consensus such as global warming and the safety and efficacy of vaccines (Weber and Stern 2011; Jacobson, Sauver, and Rutten 2015; Peretti-Watel et al. 2020). As a result, policymakers often struggle to achieve outcomes that rely on people's willingness to adjust their behaviors based on the acceptance of such scientific facts. The recent COVID-19 pandemic has brought these challenges into sharp focus: Despite an abundance of high-quality public information about the virus,[1] beliefs about its risks varied widely across individuals, affecting their willingness to follow public health guidance and engage in social distancing behaviors to reduce their risk of exposure.

But why did people with similar exposure to information from public health experts hold such divergent beliefs about the risks from COVID-19? In this project, we explore the role of individuals' social networks—their friends, families, and acquaintances—in shaping beliefs and behaviors during the COVID-19 pandemic. We first analyze the effects of friend exposure to COVID-19 cases on individuals' social distancing behavior. We document that individuals who have friends in locations with more severe outbreaks disproportionately reduce their mobility. We then study the mechanisms underlying this effect, showing that friend exposure to COVID-19 increases individuals' willingness to reduce mobility at least in part by influencing their beliefs about COVID-19. As such, information acquired through social networks shifted beliefs and behaviors even when information on the same topic had been prominently communicated by domain experts. This finding has important implications for both the design of policy and the development of new models of information acquisition.

We work with de-identified data from Facebook, a large online social networking service. The data provide information on individuals' movement

[1] A poll by the Pew Research Center (2020) conducted March 10–16, 2020, found that 89% of respondents had been following news related to COVID-19 very closely or fairly closely, with only 2% saying they had been following the news not at all closely.

patterns and the location of their friends, allowing us to measure the effects of friend exposure to COVID-19 on social distancing behavior.[2] The data also include information on public posts on the platform and membership in public Facebook groups, allowing us to study individuals' perceptions of the COVID-19 pandemic. Relative to the cell phone location data used in much of the existing research on social distancing behavior, our unique ability to link individual-level data on mobility to information on demographics, social networks, and proxies for perceptions allows us to generate novel insights into the determinants of behaviors and beliefs.

We begin by documenting time-series patterns in mobility and show that—consistent with prior work—US Facebook users in our sample drastically reduced their mobility after the outbreak of the pandemic. In mid-February 2020, the probability of staying home averaged around 18% on a given day; by late March, this probability had increased to about 30%.

We then explore the role of friendship networks in shaping social distancing behavior. To illustrate our results in the raw data, we first focus on the early onset of the pandemic. We classify each individual as being either above or below the median of friend exposure within their zip code, based on the exposure of their social network to COVID-19 as of March 15, 2020, right after President Trump declared a national emergency. Prior to the pandemic, the movement patterns of the two groups look strikingly similar. In contrast, after the outbreak, users with above-median friend exposure—that is, those who have relatively more friends living in areas highly affected by the virus—were more likely to stay home compared to others in the same zip code with lower friend exposure. Quantitatively, a 1 standard deviation higher friend exposure to COVID-19 cases was associated with an 8.8% larger increase in the probability of staying home by April 2020. These differences remain large and significant when we include controls for time-varying effects of various demographics and other characteristics of an individual's social network.

A potential concern with interpreting these cross-sectional findings is that the location of individuals' friends in the United States may be associated with other factors that could impact social distancing behaviors during

[2] We observe measures of mobility only for Facebook users who consented to sharing and storing their location information. We proxy for staying at home with staying within a single level-16 Bing tile, an area of about 600 meters × 600 meters (see sec. I.B). We use Facebook friendship links as a proxy for an individual's real-world social network, and believe that it provides a high-quality measure of the peers whom an individual would interact with both online and in the offline world. Overall, Facebook users are highly representative of the US population, and friendship links largely represent real-world friends and acquaintances (Jones et al. 2013). Indeed, prior work has shown that in the United States, Facebook friendship links provide a reliable representation of real-world friendship links (e.g., Bailey et al. 2018a, 2019, 2022a, 2022b; Chetty et al. 2022a, 2022b).

the pandemic. For instance, people with friends in early hot spots such as New York City and Seattle might be more politically liberal and, as a result, independently engage in more social distancing than their neighbors. To address concerns like these, our main specification uses a dynamic approach that estimates the effects of changes in friend exposure to COVID-19 over a given month on changes in social distancing during that month as the pandemic evolves. We demonstrate that individuals with friends in the early hot spots such as Seattle disproportionately reduced their mobility in the early pandemic compared to their otherwise similar neighbors with friends in different parts of the country. But by June 2020, it was individuals with friends in the newly emerging hot spots such as Oklahoma, Texas, and Arizona who disproportionately increased their social distancing. To interpret our results as driven by unobservables rather than as evidence for a causal effect of friend experiences on social distancing, one would need to argue that in every month, individuals with friends in regions with the largest outbreaks happened to reduce their mobility for reasons other than their friend exposure. Since a plausible version of this story is difficult to tell, we conclude that higher friend exposure to COVID-19 likely induces social distancing. We also find that the effects of friend exposure to COVID-19 on mobility patterns are virtually identical for weekends and weekdays, suggesting that the reduced mobility associated with friend exposure to COVID-19 is by choice and not due to differences in individuals' ability to work from home.

We then explore the mechanisms through which social networks affect high-stakes decisions such as whether to reduce mobility during a pandemic. In our context, a direct effect could exist if individuals in current virus hot spots schedule fewer in-person social interactions with their friends. Alternatively, a preference effect might arise if those in more affected areas engage in more homebound activities such as cooking, leading their friends to become more engaged in these activities. Finally, friend experiences might affect individuals' beliefs about the benefits of social distancing by providing information about the severity of the virus in a way that particularly resonates with the individuals.

To understand the role of these possible explanations in our setting, we first show that changes in the COVID-19 exposure of friends living more than 100 miles away still have very sizeable effects on an individual's social distancing. This suggests that a large part of our results is not driven by a direct effect of friend exposure to COVID-19 limiting visits and interactions with the affected friends.

Next, we explore whether friend experiences shape behavior by affecting individuals' beliefs. We use data from public user posts and group memberships to construct a measure of individuals' stated beliefs about COVID-19 and their attitudes toward social distancing. Friend exposure to COVID-19 cases increases an individual's propensity to post about

COVID-19 and the probability that such posts voice support for restrictions on public life. Similarly, greater friend exposure to COVID-19 cases lowers the likelihood that an individual joins public Facebook groups advocating for a reopening of the economy.

It is noteworthy that we find this effect of friend experiences on individuals' beliefs and behaviors even in a context where high-quality expert information about the risks of COVID-19 and the need for social distancing was ubiquitous and intensely covered in the media. It is thus unlikely that friends conveyed content that individuals had not already received through other channels. Instead, it is more likely that the information provided by friends—even if not necessarily new per se—resonated more with individuals and thus had a large effect on their beliefs and behaviors. Our findings therefore suggest that policymakers may have more success at shifting beliefs and behavior when relevant information is conveyed by people who resonate with the relevant target communities.

Our work speaks to a large literature on how individuals form beliefs and the extent to which these beliefs translate into actions (e.g., Malmendier and Nagel 2011; Armantier et al. 2015; Bachmann, Berg, and Sims 2015; Armona, Fuster, and Zafar 2019; Kuchler and Zafar 2019; Roth and Wohlfart 2020; Giglio et al. 2021a, 2021b; Rothwell et al. 2021; Bakkensen and Barrage 2022; Bordalo et al. 2022; D'Acunto et al. 2023). Most closely related is work that documents the possible role of social interactions—and in particular the experience of friends—on belief formation and behavior. In this literature, Bailey et al. (2018a, 2019) show that friends' house price experience can influence a person's own house price expectations. Similarly, Ratnadiwakara (2021), Hu (2022), Mayer (2023), and Xu and Box-Couillard (2023) use county-level social network data from Bailey et al. (2018a) to conclude that when an individual's friends experience extreme weather events such as hurricanes and floods, this can affect a person's own beliefs about climate change. Relative to this literature, our work uses individual-level data on social networks to highlight that friend experiences shape beliefs, opinions, and behaviors even in settings where high-quality information from domain experts is ubiquitous. This suggests that the role of friends in shaping beliefs and behaviors goes beyond those friends being a low-cost source of information, as in Banerjee et al. (2019). Instead, the evidence provides support for models of learning in which the identity of the person conveying the information matters for how much weight the information receives in the belief-formation process. Malmendier and Veldkamp (2022) propose such a model in which "abstractly learned statistics and other information tends to be weighted significantly less than information gathered from . . . the experiences of others whom we care about, identify with or empathize with."

This paper also contributes to a growing literature on the determinants of social distancing during the COVID-19 pandemic, surveyed by Rasul and Giuliano (2020) and Brodeur et al. (2021),[3] as well as work on the effect of social networks on health behaviors more generally (see Christakis and Fowler 2007, 2008; Huang et al. 2014; Fletcher and Ross 2018; Sato and Takasaki 2019). In related work, Tian, Caballero, and Kovak (2022) argue that international migration networks helped to convey information about the disease. Similarly, Charoenwong, Kwan, and Pursiainen (2020) use county-level social network data from Bailey et al. (2018b) to show that individuals living in US counties with more connections to China and Italy—two early hot spots of the COVID-19 pandemic—reduce their mobility more. Makridis and Wang (2020) show that consumption decreases more in counties with higher friend exposure to COVID-19 cases. Relative to this work, our individual-level analysis allows us to absorb any direct effects of local conditions likely correlated with friend exposure (see Kuchler, Russel, and Stroebel 2022) and our data on posts and group memberships allow us to establish individuals' beliefs about COVID-19 as an important mechanism through which friend exposure affects mobility.

I. Data and Descriptive Statistics

We work with de-identified data from the global online social networking site Facebook to measure individual-level social networks and social distancing behavior.[4] As of December 2019, Facebook had 248 million monthly active users and 190 million daily active users in the United States and Canada (Facebook 2020). Greenwood, Perrin, and Duggan (2016) found that, among US adults, usage rates were relatively constant across income groups, education levels, and race; usage rates were slightly declining in age.

Establishing a connection on Facebook requires the consent of both individuals, and a person can have at most 5,000 connections. As a result, Facebook connections are primarily between real-world friends, acquaintances, and family members and Facebook networks resemble real-world social networks more closely than networks on other online platforms where unidirectional links to non-acquaintances such as celebrities are common. Indeed, prior studies show that Facebook networks predict

[3] In this literature, civic capital (Guiso, Gulino, and Durante 2020; Barrios et al. 2021; Chetty et al. 2022a), trust in scientific knowledge (Brzezinski et al. 2021), trust in policymakers (Bargain and Aminjonov 2020), general trust (Brodeur, Grigoryeva, and Kattan 2021), news consumption (Simonov et al. 2022; Bursztyn et al. 2023), political affiliation (Allcott et al. 2020b; Barrios and Hochberg 2021), policy decisions (Allcott et al. 2020a), and potential spillover effects of policy across states (Holtz et al. 2020) have all been shown to affect social distancing.

[4] We cannot publicly share the individual-level data described in this section, but we provide the code used in our analyses in a replication package (Bailey et al. 2023).

many important real-world economic and social interactions, including patterns of trade (Bailey et al. 2021), patent citations (Bailey et al. 2018b), travel flows (Bailey et al. 2020a, 2020b), housing choices (Bailey et al. 2018a, 2019), bank lending (Rehbein and Rother 2020), social program participation (Wilson 2022), product adoption decisions (Bailey et al. 2022b), investment decisions (Kuchler et al. 2022), disease transmission (Kuchler, Russel, and Stroebel 2022), and upward income mobility (Chetty et al. 2022a, 2022b).

A. Sample Restrictions and Summary Statistics

Our analyses of mobility behavior are limited to a subpopulation of Facebook users who have consented to sharing and storing their location,[5] have active accounts, are 18 or older, live in the 50 US states or the District of Columbia, and have between 100 and 1,500 US-based Facebook friends. We restrict the analysis to ZIP Code Tabulation Areas (ZCTAs) with 50 or more users who meet all previous requirements. Overall, the sample of users who meet the above criteria includes 12.8 million individuals. The average ZCTA has 592 users, the median has 319, and the 10th percentile has 72 users. We do not require users to have location information in every week (for example, if their mobile device was turned off) and thus observe information for about 7.2 million users per week.

Table 1 provides summary statistics on the users in our mobility sample. Age ranges from 26 years at the 10th percentile to 63 years at the 90th percentile. We see that 53% of the sample is female, and just over half the users have listed a college.[6] We also observe whether a user primarily accesses Facebook from an iPhone or from an Android phone, with about 25% of the sample using an iPhone.[7] Finally, we observe that about half the sample sometimes also accesses Facebook from a tablet (e.g., an iPad).

After mapping users to their presumed ZCTA of residence, we supplement our individual-level data with public data on median household income from the 2014–18 American Community Survey (ACS). The median user in our sample lives in a ZCTA with a median household income

[5] Users consented to having their location stored if they used a feature that required high-frequency location data to function. We do not see a shift in usage patterns around the onset of the pandemic, though usage of these features had been slowly declining. To address possible concerns that the sample of users sharing their location is biased, we confirm that our core results are similar when reweighting users sharing their location so that their observable characteristics match those of our broader sample of Facebook users. We also show that our baseline patterns replicate at the zip code level using an independent source of movement data provided by SafeGraph.

[6] This measure captures college attendance better than college degree attainment, with the former much higher than the latter in the general population.

[7] All users in the mobility sample use Facebook on a smartphone, as this is required for us to observe their GPS location.

TABLE 1
Summary Characteristics: Mobility Sample

			Percentile				
	Mean	Standard Deviation	10th	25th	50th	75th	90th
Age	43.58	14.93	26	32	42	54	63
Female	.53	.50	0	0	1	1	1
Has college	.53	.50	0	0	1	1	1
Has iPhone	.25	.43	0	0	0	0	1
Has tablet	.53	.50	0	0	1	1	1
Zip code income (USD)	58,792	21,961	36,160	43,648	54,000	69,203	88,096
Number of friends	532.80	326.61	193	276	441	718	1,047
Friend exposure to cases	10.35	19.34	.74	1.77	4.49	11.12	26.31
Staying at home (February):							
All	18.33	29.35	0	0	0	28.57	66.67
Weekend	19.39	34.44	0	0	0	50.00	100.00
Weekday	16.83	29.80	0	0	0	20.00	66.67
Bing tiles visited (February):							
All	10.96	9.07	1.57	3.43	9.00	15.86	23.43
Weekend	10.57	9.79	1.00	3.00	7.50	15.50	24.50
Weekday	11.34	9.77	1.50	3.40	9.00	16.20	24.60

Note.—Summary statistics describing individuals analyzed in our mobility sample of users. Individual-level characteristics include age, gender, whether the user has a college listed on Facebook, whether the user primarily accesses Facebook mobile from an iPhone, whether the individual has accessed Facebook from a tablet, number of friends, friend exposure to COVID cases on March 15, and patterns of mobility during the week of February 25 to March 2. The table also includes information on the users' home ZCTA 2018 median household income.

of $54,000, not far from the true US median household income of $53,958. The 10th and 90th percentiles are $36,160 and $88,096, respectively, numbers that are also close to their US population equivalents of $34,658 and $89,355. For comparison, table A.1 (tables A.1–A.27 are available online) provides summary statistics for a broader population of Facebook users without the requirement for location information. This broader sample and the mobility sample are largely similar, though users in the mobility sample are slightly less likely to have attended college, less likely to use an iPhone, and are from slightly lower-income ZCTAs on average.

To measure an individual's exposure to COVID-19 cases, we use data from Dong, Du, and Gardner (2020) on COVID-19 cases at the county-by-day level. We map each user to a county by crosswalking their ZCTA of residence to the county in which the largest fraction of the ZCTA's population resides.[8]

[8] We use COVID-19 cases rather than deaths. The average friend exposure in our sample in March is 10 cases as shown in table 1, so regressions using death instead of cases would be underpowered given the still relatively low mortality rates.

B. Measuring Mobility and Social Distancing

We measure mobility using user-level GPS data for individuals who have consented to sharing and storing their location information. These location data are recorded at high frequency: In Iyer et al. (2023), researchers noted that 54% of users globally who opted into this feature record a location "ping" in at least half of the 5-minute intervals during each day.[9] Location data are aggregated using the Bing Maps Tile System, which defines a series of grids at different resolution levels over a rectangular projection of the world (Schwartz 2018). We use level-16 Bing tiles, which are 600 meters $\times$ 600 meters at the equator. We construct two mobility indexes: (i) whether a user remains in the same level-16 Bing tile throughout the day (which we will refer to as "staying at home"), and (ii) the total number of distinct level-16 Bing tiles visited on a given day.

Figure 1 shows daily values of our two mobility measures between early February and late May 2020.[10] In figure 1*A*, we see that in February and early March, between 15% and 20% of users stayed at home on a given day, with recurring spikes on weekends (see also table 1). Starting the week of March 16—the first week after COVID-19 had been declared a national emergency and when a large number of schools and offices were closed in response to the emerging pandemic—the probability of staying at home jumped to well over 30% by March 23. It rarely fell below 30% throughout April. In May, as social distancing restrictions were eased across parts of the United States, the series decreased steadily, though the probability of staying at home remained elevated relative to the baseline period and never fell below 20%. Figure 1*B* shows that the average number of tiles visited follows the same patterns over time. Thus, in our main analysis, we focus on the probability that a user stays at home as our primary mobility metric.

We also briefly explore how the extent of social distancing varies with individual characteristics—something our individual-level data are uniquely suited to examine. Tables A.2 and A.3 show that while older individuals already spent more time at home prior to the pandemic, they

[9] These data are similar to those described in Maas et al. (2019) and used to create the Facebook Data for Good Mobility Dashboard, available at https://www.Covid19mobility.org/dashboards/facebook-data-for-good.

[10] In all graphs in this section, we control for the possible effects of a technical change in the methodology of location data collection near the end of February. Specifically, we assume that the relationship between the levels of our metrics in early February and the levels in the week of February 24 matches the relationship over the same time periods in the SafeGraph data described in the appendix, available online. Such an adjustment is not necessary in any other analysis in the paper, where we either use only data after the technical methodology change or estimate results using a difference-in-differences approach (where the methodology change had quasi-random effects across groups).

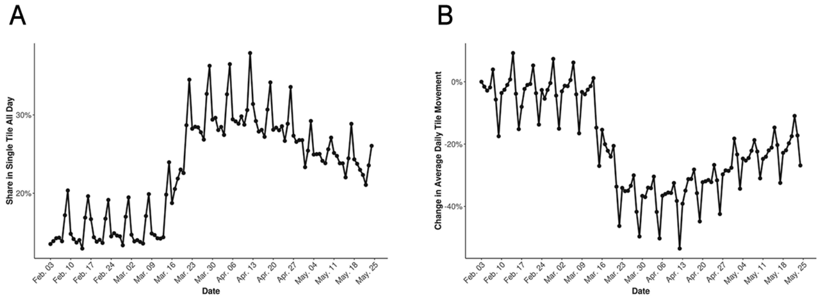

FIG. 1.—Mobility over time. The panels show average mobility patterns according to two metrics described in section I.B. *A*, Probability of staying at home. *B*, Percent change in average number of tiles visited from February 3.

changed their behavior more during the pandemic, consistent with the fact that COVID-19 poses a greater risk to that demographic. Similarly, female users increased their rate of staying home by 4.5 percentage points more than men did, consistent with an increased childcare burden being borne by women during the pandemic (see Alekseev et al. 2022). We also find that users who list a college education increased their probability of staying home by more than users without college education. This finding is consistent with the conclusions from Dingel and Neiman (2020), who note that jobs requiring high levels of educational attainment are less likely to be deemed "essential" and can more often be done from home (though we find that individuals listing a college degree were also more likely to stay at home on weekends).[11]

II. Effects of Friend Exposure to COVID-19 Cases on Social Distancing

We next explore the relationship between friend exposure to COVID-19 cases and social distancing behavior. We first study behavior at the onset of the pandemic, allowing us to illustrate our results in the raw data. In our primary specification, we estimate the effect of changes in friend exposure on changes in social distancing as the pandemic progresses, allowing us to rule out possible concerns about persistent unobservable differences correlated with friend exposure to COVID-19.

[11] We present time-series versions of these results in figs. A.1 and A.2 (figs. A.1–A.15 are available online). These figures highlight that the demographic differences in social distancing behavior discussed above arise in mid-March 2020 and persist through the end of May.

A. Friend Exposure and Social Distancing Behavior at Onset of COVID-19 Pandemic

We measure friend exposure to COVID-19 cases at the onset of the pandemic for each user as

$$\text{FriendExposure}_i^{\text{Mar15}} = \sum_{j=1}^{J} \text{FracFriends}_{ij}^{\text{Mar15}} \times \text{Covid19Cases}_j^{\text{Mar15}}. \qquad (1)$$

The term $\text{FracFriends}_{ij}^{\text{Mar15}}$ is the share of US-based friends of person i in county j on March 15, and $\text{Covid19Cases}_j^{\text{Mar15}}$ is the cumulative number of COVID-19 cases reported in county j before March 15. The mean number of cases across counties is 0.95 and the standard deviation is 9.33.[12]

Table 1 shows substantial variation in this measure of friend exposure across individuals, with a mean of 10.4 friend-weighted cases and a standard deviation of 19.3. For the first few weeks of the pandemic, the correlation of FriendExposure_i measured at different points in time is high, as similar US locations had the highest cumulative case counts. This finding also suggests that strategic friendship formation after the discovery of COVID-19 does not drive our results (see fig. A.4 for details).

1. Friend Exposure and Social Distancing Behavior at the Onset of COVID-19 Pandemic—Raw Data

We first focus on users within the same ZCTA and compare the social distancing behavior of those with high and low levels of friend exposure. Concretely, for every ZCTA, we calculate the median friend exposure to COVID-19 cases as of March 15. We then define HighExp_i for user i as an indicator of whether their friend exposure is higher or lower than the median in their home ZCTA. This measure of relative exposure allows us to show variation in social distancing by friend exposure in the raw data.

Figure 2*A* presents a time-series plot for the probability of staying at home split by HighExp_i. Before the onset of the pandemic, there are no differences in movement patterns between users in the same ZCTA with high and low levels of friend exposure. In February the probability

[12] Figure A.3 maps the number of cases by county. In this section, we primarily use measures of COVID-19 cases that do not normalize cases by the county populations. In the early stages of the pandemic, when measured case counts were low, the raw number of cases was likely a more salient measure of COVID-19 exposure than a normalized measure. For example, the areas with highest case exposures on March 15 were King County and New York City, each widely covered as early pandemic hot spots. By contrast, the areas with highest per capita infection rates were Pitkin and Eagle counties in Colorado. The outbreaks in these small counties received relatively little attention. In col. 3 of table A.9 we show that our primary results hold when normalizing case counts by population. In sec. II.B we use normalized measures of exposure when exploring later stages of the pandemic.

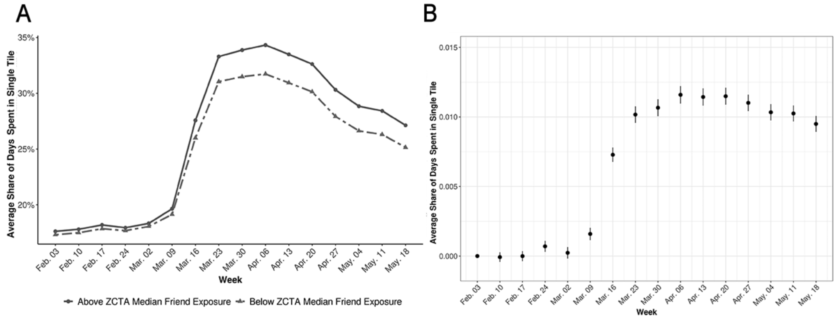

FIG. 2.—Effects of friend exposure to COVID-19 on probability of staying at home. Panels show the relationship between friend exposure to COVID-19 on March 15 and mobility behavior. We measure the latter as the weekly averages of the probability of staying at home from the week of February 3 to the week of May 18, separately for individuals above and below the median level of friend exposure in their ZCTA. *A*, Raw means. *B*, Coefficients estimated using the difference-in-differences setup specified in equation (3). The specification includes fixed effects for each individual, and fixed effects for the following groups, interacted with dummies for each week: ZCTA, age group, gender, has college listed on Facebook, has iPhone, has tablet, and percentiles of friend exposures (as in eq. [3]) for median household income, population density, and the share of the population living in urban areas. Standard errors are clustered by ZCTA. See figure A.5 for a corresponding analysis of the average number of tiles visited.

of staying at home for both groups was between 17% and 20%, with any differences always less than half of a percentage point. Starting in mid-March, however, users with high friend exposure to COVID-19 became substantially more likely to stay home. By early April, individuals with high friend exposure have a probability of staying at home of close to 35%, compared to less than 32% for users with lower levels of friend exposure.

2. Difference-in-Differences Analysis

While the raw data show identical mobility patterns between individuals with high and low friend exposure to COVID-19 prior to the pandemic, both in levels and in changes, it is important to acknowledge that friend exposure is likely nonrandom even within a ZCTA: Given the geographic concentration of US COVID-19 cases in mid-March, friend exposure likely correlates with individual characteristics that might also affect behavior during a pandemic (but not before).[13] We next show the importance of controlling for such observable differences before introducing our main

[13] We present summary statistics of the high- and low-exposure samples in table A.4. To understand the relationship between friend exposure to COVID-19 and individual and ZCTA-level characteristics, we regress a set of control variables on the log of $\text{FriendExposure}_i^{\text{Mar15}}$ in table A.5. We find that certain demographics are indeed correlated with friend exposure on March 15. For example, older users and those reporting college attendance had higher levels of friend exposure.

specification, which uses a dynamic approach to also address concerns about unobservable factors. Figure 2*B* shows estimates of β_t from the following difference-in-differences specification:[14]

$$Y_{it} = \mu_i + \sum_{t=1}^{15} \beta_t (\text{HighExp}_i \times \text{week}_t) + \sum_{t=1}^{15} \delta'_t (X_i \times \text{week}_t) + \varepsilon_{it}. \quad (2)$$

Here Y_{it} is individual i's mobility during week t. We include data for the week of February 3 as $t = 0$, but omit a coefficient for this reference time period. The term μ_i is an individual-level fixed effect, and HighExp_i is an indicator equal to 1 if user i has friend exposure greater than their ZCTA median on March 15; week_t is an indicator for the week of the outcome. The vector X_i includes fixed effects for the individual's location (ZCTA), college attendance, ownership of iPhone and tablet, age group, and gender. It also includes fixed effects for percentiles of friend-weighted median household income, population density, and share urban, each calculated analogously to our friend-based COVID-19 exposure as[15]

$$\text{FriendMetric}_i = \sum_{j=1}^{J} \text{FracFriends}_{ij} \times \text{Metric}_j. \quad (3)$$

Relative to the simple comparison of means in figure 2*A*, figure 2*B* allows for time-varying differences across individuals with different demographics and different distributions of friendship networks across measures such as the average income or population density where friends live. Consistent with figure 2*A*, the two groups' movements look nearly identical prior to the pandemic. Users with higher friend exposure are substantially less mobile after the outbreak begins, though the inclusion of the rich set of control variables in equation (2) somewhat reduces the estimated magnitude of the difference.[16]

Finally, to benchmark the magnitude of this effect, we use a multivariate analysis to compare the relative magnitudes of changes in friend exposure to COVID-19 on social distancing against the differences in social distancing across demographic groups (see app. A.1). We find that

[14] Since "treatment timing" does not vary, this simplifies our specification relative to that estimated in Goodman-Bacon (2021).

[15] The data on median household income and population density come from the 5-year ACS from 2014–2018 and the share of the population living in urban areas comes from the 2010 census.

[16] In figs. A.6 and A.7, we estimate eq. (2) separately for weekdays and weekends. We find that individuals with high friend exposure tend to reduce their mobility by a similar amount on both weekends and weekdays, which is consistent with a mechanism in which voluntary social distancing drives our results, as opposed to mechanisms related to one's industry of employment or ability to work from home. These two figures also show specifications that include college-by-week fixed effects (i.e., a week-specific mobility effect for everyone who attended the University of Michigan), further demonstrating the robustness of our results.

a 1 standard deviation increase in friend exposure to COVID-19 corresponds to an increase in social distancing that is more than two-thirds as large as the effect of being age 55 or older (relative to being below age 35), and roughly half of the effect of reporting a college.

B. *Dynamics of Friend Exposure to COVID-19 and Social Distancing Behavior over Time*

We now turn to our primary specification to estimate the effect of friend exposure to COVID-19 cases on social distancing behavior. Rather than focusing on the effects of friend exposure at the onset of the pandemic, we now study the effects of changes in friend exposure as the pandemic evolves on changes in social distancing. As the pandemic progressed, the changing geography of COVID-19 outbreaks led different individuals to experience increases in friend exposure at different points in time. With fixed individual characteristics differenced out, the dynamic approach therefore alleviates important concerns that correlations between our friend-exposure measure and unobservable individual characteristics could be driving our earlier results.

1. Measuring Changes in Friend Exposure to COVID-19

For each month, we define changes in an individual's friend exposure to COVID-19 cases as follows:[17]

$$\begin{aligned}\text{ChangeFriendExposure}_{it} = {} & \log(1 + \text{FriendExposure100}k_{it}) \\ & - \log(1 + \text{FriendExposure100}k_{it-1})\end{aligned} \tag{4}$$

with $\text{FriendExposure100}k_{it} = \Sigma_{j=1}^{J}\text{FracFriends}_{ijt}(\text{Covid19Cases}_{jt}/\text{Residents100}k_j)$. Figure 3 shows the locations with the largest changes in per capita COVID-19 cases in each month in the sample, with brighter shades corresponding to larger increases. In March, case growth was highest in New York, Seattle, Denver, and Louisiana. By April, the highest case growth was in the Midwest; in May, hot spots appear in Minnesota, Iowa, and North Carolina, while in June the location of hot spots moved to Texas, Oklahoma, and Arizona. In July, southern Texas and the northwestern

[17] In our dynamic analysis, we normalize cases by population, since, as the pandemic progressed, coverage of hot spots shifted from talking about "total cases" to "total cases per population"; see also footnote 12. Using the difference of logs gives a higher weight to the same absolute increase of cases per population in places with relatively fewer prior cases per population. While we believe that this is a useful specification to capture salient changes in COVID-19 exposure, we have verified that our results and conclusions are robust to a wide variety of ways of measuring changes in friend exposure to COVID-19.

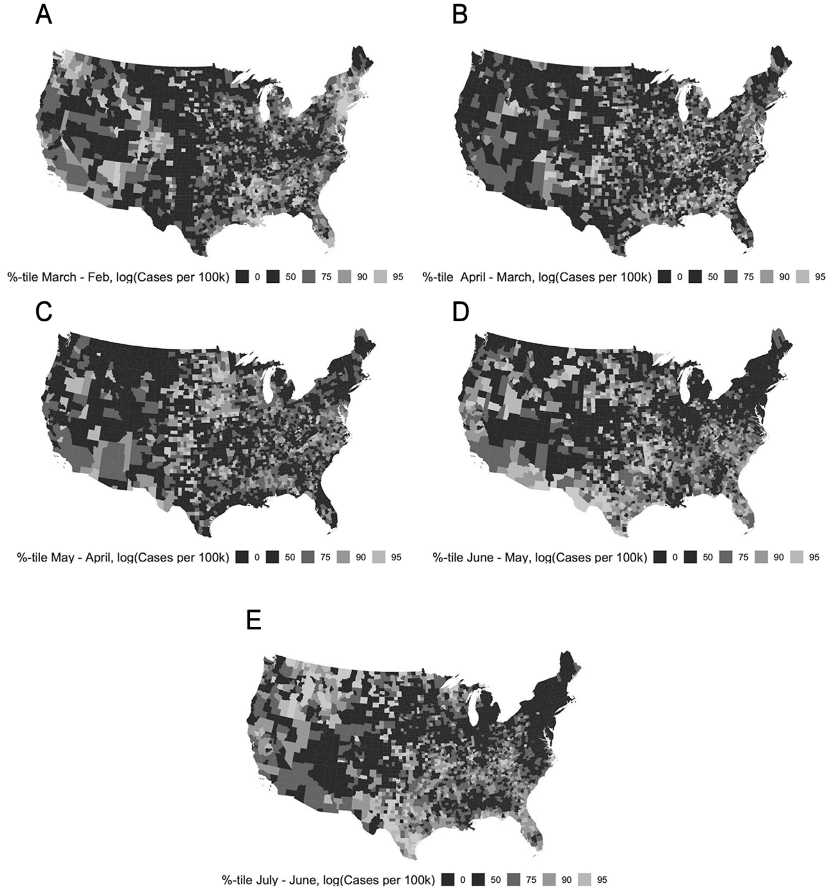

FIG. 3.—Variation in Δ COVID-19 cases per capita. Panels show percentiles of the change in log(COVID-19 cases per 100,000 residents + 1) by county for the continental United States. Cases are measured on the last Friday of each month. *A*, Change from February to March; *B*, change from March to April; *C*, change from April to May; *D*, change from May to June; *E*, change from June to July. Darker shade indicates a smaller increase and lighter shades indicate a larger increase.

mountain states see new hot spots emerge. This geographic variation in case growth throughout our sample means that, in each month, it is different individuals who happen to be most exposed to COVID-19 case growth through their friendship networks. Indeed, table A.6 shows that the correlation between changes in friend exposure to COVID-19 and demographic characteristics changes over time. For example, individuals with a listed college were more exposed to COVID-19 case growth through their friends at the beginning of the pandemic; in later months, as the pandemic spread across the United States, this relationship reversed.

2. Effect of Changes in Friend Exposure to COVID-19 on Changes in Social Distancing Behavior

To analyze the effects of changes in friend exposure over time on changes in social distancing behavior, we estimate the following equation:

$$\Delta Y_{i,t} = \sigma_0 + \sigma_1 \text{ChangeFriendExposure}_{i,t} + \sigma_{2,t} X_{i,t} + \epsilon_{i,t}. \tag{5}$$

Here $X_{i,t}$ captures a range of characteristics of individual i at time t. In our baseline specification, $X_{i,t}$ includes fully interacted month × ZCTA × age group × gender × has college × has tablet × has iPhone fixed effects. This interaction captures any changes in (or varying effects of) local conditions and lets their effects covary with characteristics. We also include percentiles of friend-weighted urbanity, population density, and median household income as defined in equation (3), each interacted with month fixed effects to allow the effects of those network characteristics on changes in social distancing to vary over time.[18]

Column 1 of table 2 presents the estimate of σ_1 from equation (5), pooling across all months in our sample. Figure A.10 presents the corresponding binned scatter plot. The results show that doubling the increase in friend exposure is associated with a 9% higher change in the likelihood that a person stays at home in a given month.

We also explore the relationship between changes in social distancing behavior and changes in friend exposure to COVID-19 for each month separately. This allows us to explore whether the effects in the pooled regression in column 1 were primarily driven by individuals' social distancing behavior in a given month. Concretely, we estimate equation (5) separately for each month and include all past changes in friend exposure as explanatory variables. Columns 2–6 of table 2 present the results of this analysis. In March—consistent with our earlier results for the onset of the pandemic—higher increases in friend exposure significantly increase the probability of staying at home. Importantly, in subsequent months, changes in social distancing behavior are driven primarily by changes in friend exposure in the corresponding months. That is, individuals with friends in early hot spots, such as New York City, Seattle, Denver, and Louisiana, stay home more in March than their otherwise similar neighbors with friends in the Midwest. As the pandemic progresses, their neighbors with friends in the Midwest experience large increases in friend exposure in April and, accordingly, increase their probability of staying at home. In May, hot spots appear in Minnesota, Iowa, and North Carolina, and again, individuals with friends in those areas start social

[18] Note that these measures are calculated using the friend network as of March 15. The network is relatively constant over the sample period and recalculating the measure using alternative exposure dates does not change our results.

TABLE 2
EFFECTS OF FRIEND EXPOSURE BY MONTH: Δ PROBABILITY OF STAYING AT HOME

	Monthly Change in Probability of Staying at Home					
	All Months (1)	March (2)	April (3)	May (4)	June (5)	July (6)
Change friend exposure:						
Same month	.208***					
	(.029)					
March		.207***	.006	−.076**	.097	.037
		(.046)	(.040)	(.048)	(.054)	(.064)
April			.035	.096	.329***	.069**
			(.052)	(.056)	(.061)	(.071)
May				.379***	.044	−.057
				(.082)	(.078)	(.094)
June					.854***	−.329*
					(.114)	(.127)
July						.323**
						(.138)
Other network exposure fixed effects	Yes × month	Yes	Yes	Yes	Yes	Yes
Zip code × age group × gender × has college × has tablet × has iPhone	Yes × month	Yes	Yes	Yes	Yes	Yes
R^2	.211	.174	.141	.150	.146	.145
Sample mean	1.611	14.214	−.923	−5.989	−1.068	.679
Observations	30,742,008	6,688,448	6,579,359	6,169,176	5,848,722	5,456,303

NOTE.—Column 1 reports the results of regression (5) with one observation per user per month between March 2020 and July 2020. Change in friend exposure is defined in eq. (4). In cols. 2–6, we subset the data to observations from the months up to and including the one listed in the header. Here, each observation is an individual. In all columns, the outcome variable is the change in the probability of staying home between the final week of a given month and the final week of the previous month. We define the final weeks to be the last Friday to Thursday period in a month. The last weeks are then February 25–March 2, March 24–March 30, April 21–April 27, May 26–June 1, June 23–June 29, and July 21–July 28. The sample of users is restricted to those for whom location can be observed at the end of each of the two relevant months. In all columns we control for interactions of ZCTA fixed effects, age groups, gender, whether the individual has a college listed on Facebook, whether the individual primarily accesses mobile Facebook from an iPhone, and whether the individual has accessed Facebook from a tablet. All columns also include fixed effects for percentiles of friend exposures (as described in eq. [3]) for median household income, population density, and share of the population living in urban areas. Standard errors are clustered by ZCTA.
* $p < .10$.
** $p < .05$.
*** $p < .01$.

distancing more than their otherwise similar neighbors with friends in other parts of the country. Across all months, we find that the most recent changes in the rate of friend exposure are the most important, though our results in April are not statistically significant. These findings support our hypothesis that friend exposure to COVID-19 has a sizeable effect on social distancing behavior.

As shown above, characteristics of users with high friend exposure to changes in COVID-19 cases vary substantially over time. As a result, the dynamic relationship between changes in friend exposure and changes in social distancing behavior allows us to establish this relationship without the concern for bias from unobservable characteristics and to overcome several of the shortcomings of the cross-sectional specification studied before. In particular, any individual characteristics with a constant effect on the level of mobility are differenced out.[19] In addition, the effects of observable characteristics on social distancing are controlled for, even if the relationship between characteristics and the change in social distancing behavior varies across months. Similarly, the specification allows the effects of all local conditions on social distancing to vary by characteristics and over time. To obtain unbiased estimates of the causal effect of friend exposure on social distancing with specification (5), we need to assume that any time-varying effect of unobservable characteristics on social distancing is not systematically correlated with the changes in friend exposure—a very plausible assumption.

We conduct several robustness checks to the analysis presented in table 2. Table A.10 shows that our results are very similar when focusing only on users for whom a complete panel is available. Similarly, including an individual fixed effect to capture possible individual-level trends in mobility over time does not affect the results. Table A.11 shows that this relationship holds when using the number of tiles visited as the outcome and when using a Poisson functional form. In table A.12, we show that we obtain similar patterns when regressing changes in mobility only on changes in friend exposure for the same month (without also including changes in prior months).

C. *Heterogeneity of Friend-Exposure Effects*

We next explore heterogeneity in the effect of friend exposure on social distancing behavior along an individual's own characteristics. To avoid capturing heterogeneity in the ability to work from home rather than the desire to stay home, we focus on weekend movements. Specifically, we modify equation (5) to interact our measure of changes in friend exposure with indicators for various demographic characteristics. Table 3 shows that changes in friend exposure have a larger effect on the social

[19] Instead of an individual mobility effect, one could assume an individual-level social distancing effect that only affects mobility after the onset of the pandemic. To difference out such an effect, we can exclude the first month of prepandemic data in the estimation. Table A.10 shows that the results are very similar.

distancing behavior of younger users: the effect for those aged 35–55 is only about one-third the size of the effect for those aged 18–34. The effects of friend exposure on mobility is also substantially larger for females than for males and, similarly, larger for users with a listed college than for users without a listed college. In addition, the effect is increasing in the average income of an area, as well as the prevalence of COVID-19 in the user's own county. Interestingly, despite these heterogeneities, for nearly all of the various groups we consider, we find that increases in friend exposure lead to increases in social distancing.

We also study heterogeneity in the effects of friend exposure to COVID-19 by the strength of the underlying friendships. Friendships are ranked by "closeness" based on the extent of various interactions between users on Facebook. Our specification amends equation (5) by replacing $\text{ChangeFriendExposure}_{it}$ among all friends with four variables that measure changes in friend exposure among friends with different friend ranks: 1–25, 26–50, 51–75, and 76–100. Column 6 of table 3 shows that the effects of friend exposure tend to be strongest for the closest friends, with effect size falling off among more marginal friends. The effect of friend exposure among a person's 25 closest friends is nearly 3 times stronger than the effect of friend exposure among the person's next 25 closest friends (those ranked 26–50). The effect size is smaller, and no longer significant at 5%, for the two more distant friend groups. The decrease in the effect size of friend exposure as we move toward more distant friends is consistent with our hypothesis that the observed effects on health behavior are indeed driven by friend exposure to COVID-19 cases rather than omitted variables.

III. Mechanisms

So far we have shown that friend exposure to COVID-19 cases induces individuals to engage in more social distancing. In this section, we explore possible mechanisms behind these findings. First, there might be a direct effect of friend exposure on one's own movement, for example if individuals cut back on meeting up with friends in areas with high COVID-19 caseloads. Second, it is possible that the effect operates through a preference channel. This could occur, for instance, if homebound friends in highly exposed areas begin to bake or garden, and share tips that make these activities more appealing relative to alternative activities taking place outside one's home. Finally, friend exposure might change people's beliefs or attitudes toward the risks of COVID-19. We combine several pieces of information from users' activity on Facebook to conclude that a key part of the mechanism through which friend exposure to COVID-19 affects social distancing is through influencing individuals' beliefs about COVID-19.

TABLE 3

HETEROGENEITY OF MONTHLY FRIEND-EXPOSURE EFFECTS, WEEKENDS

	Monthly Change in Probability of Staying at Home, Weekends						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
(1) Change friend exposure × $I(\text{age} < 35)$	.745***						
	(.069)						
(2) Change friend exposure × $I(\text{age } 35\text{–}55)$	.229***						
	(.055)						
(3) Change friend exposure × $I(\text{age} > 55)$	.066						
	(.072)						
(4) Change friend exposure × female		.505***					
		(.055)					
(5) Change friend exposure × male		.109**					
		(.054)					
(6) Change friend exposure × college			.516***				
			(.054)				
(7) Change friend exposure × no college			.101*				
			(.055)				
(8) Change friend exposure × zip income first tertile				.016			
				(.064)			
(9) Change friend exposure × zip income second tertile				.252***			
				(.063)			
(10) Change friend exposure × zip income third tertile				.776***			
				(.072)			
(11) Change friend exposure × county cases first tertile					.099*		
					(.055)		
(12) Change friend exposure × county cases second tertile					.415***		
					(.079)		
(13) Change friend exposure × county cases third tertile					.687***		
					(.074)		
(14) Change friend exposure, friends ranked 1–25						.289***	
						(.039)	

(15) Change friend exposure, friends ranked 26–50						.105***	
						(.041)	
(16) Change friend exposure, friends ranked 51–75						.016	
						(.040)	
(17) Change friend exposure, friends ranked 76–100						−.067*	
						(.039)	
(18) Change friend exposure, friends < 100 miles away							.387***
							(.085)
(19) Change friend exposure, friends > 100 mile away							.321***
							(.058)
(20) Other network exposure fixed effects	Yes × month	Yes × month	Yes × month	Yes × month	Yes × month	Yes × month	Yes × month
(21) Zip code × age group × gender × has college × has tablet × has iPhone	Yes × month	Yes × month	Yes × month	Yes × month	Yes × month	Yes × month	Yes × month
(22) R^2	.189	.189	.189	.189	.189	.189	.246
(23) Sample mean	1.436	1.436	1.436	1.436	1.436	1.436	1.615
(24) F-test (rank 1–25 = rank 76–100)						34.593***	
(25) Observations	27,821,521	27,821,521	27,821,521	27,821,521	27,821,521	27,797,612	10,656,616

NOTE.—Results from versions of regression (5). The variable ChangeFriendExposure$_{it}$ is interacted with age groups in rows 1–3, gender in rows 4 and 5, whether the individual has a college listed in Facebook in rows 6 and 7, ZCTA median household income in rows 8–10, and county-level cases of COVID-19 in rows 11–13. Rows 14–17 show measures of change in friend exposure constructed using friends of certain ranks (i.e., a measure for how close friends are). Rows 18 and 19 show measures of change in friend exposure constructed using friends who live within (outside) 100 miles. Column 7 is restricted to users who have at least 100 friends <100 miles away and >100 miles away. All columns include controls for percentiles of friend exposures (as described in eq. [3]) for median household income, population density, and the share of the population living in urban areas, each interacted with month. All columns include monthly fixed effects for every group constructed from interacting ZCTA, age group, gender, has college, has tablet, and has iPhone. Standard errors are clustered by ZCTA. In table A.13 we repeat these analyses studying mobility on all days, not just on weekends. In tables A.15 and A.16 we conduct related heterogeneity analyses using versions of the specification in eq. (A.1) (available online) that focuses on exposure at the onset of the pandemic. In figs. A.8 and A.9 we show heterogeneity results in the event-study framework described in sec. II.A.

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

A. *Direct Effects of Friend Exposure to COVID-19*

We first consider the possibility that the effects of friend exposure operate largely through a direct channel, where higher COVID-19 rates in friends' locations directly reduce visits to and movements together with these friends. To do this, we perform variants of regression (5), splitting friends into two groups, one farther and one closer than 100 miles from the user's ZCTA. For this analysis, we restrict ourselves to users who have at least 100 friends within 100 miles of their ZCTA and 100 friends more than 100 miles away. Column 7 of table 3 shows that higher exposure to COVID-19 cases among all types of friends is associated with a higher likelihood that the user stays home on a given day. In addition, the impact of faraway friends relative to nearby ones is only slightly smaller in magnitude. Since trips to visit faraway friends are uncommon, our finding that COVID-19 cases in the locations of these friends have a substantial effect on an individual's mobility patterns suggests that the effects we observe are not primarily explained by a decreased likelihood of travel to visit friends in affected areas.[20]

B. *The Role of Beliefs*

We next explore whether friend exposure to COVID-19 cases affects social distancing behavior through shaping beliefs about the risks from COVID-19. To do this, we examine whether proxies for individuals' beliefs react to friend exposure to COVID-19.

1. Posting Behavior

We begin by analyzing users' public Facebook posts, which can be viewed by any other user on the platform. We use these public posts to construct two measures. First, we use regular expression searches to measure the percentage of a user's public posts that mention the coronavirus; this measure captures the user's level of general engagement in discussions about COVID-19. Second, we identify common phrases used to support or oppose social distancing measures to quantify a user's level of opposition to these measures. Specifically, we measure the number of posts opposed to social distancing as a fraction of all "signed" posts, that is, all

[20] Friends who live farther away are generally less close. At the same time, users with substantial numbers of faraway friends may have fewer close local friends, e.g., because they only recently moved to the area. To address these concerns, we additionally divide the groups according to the ranking of friend strength used in table 3, allowing us to compare friends who are similarly socially close but live different distances away from the user. The results of these regressions are presented in table A.14 and support the notion that faraway friends have substantial effects on social distancing behavior.

posts identified as either supporting or opposing these measures. Appendix C provides details on these classifications.

We estimate the effect of friend exposure to COVID-19, as well as other individual- and ZCTA-level characteristics, on these public posting behavior outcomes using the following regression:

$$Y_i = \delta_0 + \delta_1 \log(\text{FriendExposure}_i^{\text{Mar15}}) + \delta_2 X_i + \epsilon_i. \quad (6)$$

Here Y_i corresponds to one of the posting outcomes described above, and $\text{FriendExposure}_i^{\text{Mar15}}$ is defined as in equation (1). We control for fully interacted ZCTA × age group × gender × has college × has table × has iPhone fixed effects. For this analysis of users' beliefs about COVID-19, we require that users have posted publicly at least once in February, March, or April of 2020. Since we do not limit the sample to users with location sharing and storage permissions, our sample size increases substantially compared with the prior analysis.[21] Summary statistics for this sample are shown in table A.18.

Table 4 presents estimates of the coefficient of interest, δ_1.[22] In column 1, we explore the effect of friend exposure to COVID-19 on the share of public posts between February and April 2020 that are about the coronavirus. Friend exposure to COVID-19 cases has substantial effects on posting behavior: a doubling in friend exposure corresponds to an increase in the share of posts about the coronavirus of about 0.17 percentage points, a 10% increase relative to the average, even with our tight controls for ZCTA interacted with individual characteristics.[23]

This first analysis suggests that users with higher levels of friend exposure to COVID-19 are generally more likely to talk about the coronavirus, but does not capture the nature of individuals' posts. Specifically, our measure includes both posts supportive of the notion that the virus poses a great threat to public health and endorsing measures to contain the risk, and posts that downplay the threat of the virus or that call for an end to restrictions. In column 2 of table 4, we thus explore the share of

[21] We still observe an assumed ZCTA of residence based on IP address, profile information, and other factors, allowing us to include ZCTA-level controls in our regressions.

[22] In table A.17 we also measure the general sentiment of public posts relating to COVID-19 using the VADER algorithm described in Hutto and Gilbert (2014). We replace Y_i in eq. (6) with the change in average post sentiment between February 3–23 and April 6–26. We find that users with higher levels of friend exposure to COVID-19 cases have significantly larger decreases in post sentiment, suggesting the overall sentiment in their posts becomes more negative. While this result is consistent with friend exposure affecting beliefs about COVID-19 risks (e.g., as captured by posts such as, "I really hate COVID"), our measure can also pick up a wide variety of beliefs. For instance, posts critical of COVID-19–related policies (e.g., "I really hate COVID lockdowns") also display negative sentiment, complicating interpretation.

[23] Figure A.11 shows a binned scatter plot that corresponds to our analysis in col. 1. The relationship between the percentage of posts about COVID-19 and friend exposure is strong, with a functional form that is close to linear.

TABLE 4
POSTING BEHAVIOR AND GROUP MEMBERSHIP

	Dependent Variable			
	Share Posts about COVID-19 (February–April)	Share Signed Posts Opposed to Distancing (February–April)	Member Reopen Group by June 28, 2020	
	(1)	(2)	(3)	(4)
log(friend exposure March 15)	.249***	−1.929***	−.129***	
	(.006)	(.245)	(.007)	
log(friend exposure end of June)				−.122***
				(.009)
Percentiles of total number of groups (February 2020)			Yes	Yes
Other network exposure fixed effects	Yes	Yes	Yes	Yes
Zip code × age group × gender × has college × has tablet × has iPhone	Yes	Yes	Yes	Yes
Sample	People with any posts February–April	People with signed posts February–April	People with group memberships	People with group memberships
R^2	.060	.445	.074	.074
Sample mean	1.755	35.979	1.216	1.216
Observations	34,528,373	277,776	119,145,833	119,153,786

NOTE.—Results from regressions (6) and (7). Each observation is an individual. The outcome in col. 1 is the percentage of individual posts that are about COVID-19; in col. 2 it is the percentage of pro- or anti-distancing posts that are anti-distancing; in cols. 3 and 4 it is whether the individual was a member of a reopen Facebook group as of June 28. For ease of interpretation and because of small magnitudes, we rescale coefficients and standard errors by 100, so that they correspond to percentages. Post classification is based on the regular expression in app. C. Group classification is determined by the regular expression described in app. C. All columns control for percentiles of friend exposures (as described in eq. [3]) of median household income, population density, and the share of the population living in urban areas. All columns include fixed effects for every group constructed from interacting ZCTA, age group, gender, has college, has tablet, and has iPhone. The group-based analyses in cols. 3 and 4 also include fixed effects for the percentile of the number of groups an individual was in as of February 2020. Standard errors are clustered by ZCTA.

*** $p < .01$.

signed posts that oppose social distancing requirements and shutdowns. For this analysis, we concentrate on those users who share at least one signed post in February, March, or April of 2020. Friend exposure to COVID-19 decreases the likelihood that users oppose social distancing measures in their posts (fig. A.11*b* shows the corresponding binned scatter plot): a doubling in friend exposure corresponds to a 1.3 percentage point reduction in the share of signed posts opposing distancing. This implies a 4% reduction given a baseline average of 36%.[24]

2. Group Membership

We next explore the effects of friend exposure to COVID-19 cases on a user's decisions to join various Facebook groups advocating to reopen the economy. Facebook users can create and join groups to chat, meet, and otherwise engage with others. For our analysis, we focus on membership in public groups, which any Facebook user can access without additional restrictions. Since no restrictions on posting behavior or location settings are necessary for this part of the analysis, we focus on all active users who meet the nonmobility sample requirements described in section I. We present summary statistics for this group of users in table A.1.

To measure beliefs about the risks of COVID-19, we focus on groups created between March 1 and June 28, 2020, with names that suggest support for an early reopening of the economy. Appendix C provides details on how we identify these groups. We then estimate

$$\text{ReopenGroup}_i = \gamma_0 + \gamma_1 \log(\text{FriendExposure}_i) + \gamma_2 X_i + \epsilon_i, \quad (7)$$

where ReopenGroup_i is an indicator equal to 1 if, on June 28, user i is a member of at least one group advocating for the lifting of COVID-19 related restrictions. FriendExposure_i and X_i are defined as above. In addition to the control variables used in the previous specifications, we include fixed effects for percentiles of the number of groups the user is a member of as of February 2020, allowing us to control for potential differences in usage of the groups feature on Facebook. Table 4 presents estimates of γ_1 using friend exposure on March 15 (col. 3) as well as cumulative friend exposure to COVID-19 cases through June (col. 4).[25]

[24] It is possible that some of the observed effect is driven by users changing what they decide to share in the face of anticipated backlash or support from friends in hard-hit areas. However, the fact that changes in friend exposure also induce real changes in behavior that are not visible to friends in faraway locations suggests that observed effects on stated beliefs and opinions likely correspond to true changes in beliefs.

[25] In table A.17 we use a looser set of controls and also present estimates of γ_2. Table A.19 uses a normalized measure of exposure at the end of each month in our sample period. The table shows a negative, though not always statistically significant, effect in each month.

About 1.2% of all users are a member of at least one reopen group. Column 3 shows that a doubling in friend exposure to COVID-19 on March 15 decreases the probability of being a member of such a group by about 0.09 percentage points, or 7.5%. Column 4 shows that these results are similar when using cumulative friend exposure by the end of June.[26]

C. *Mechanisms: Summary and Discussion*

Taken together, the results in this section suggest that the exposure of one's friends to COVID-19 cases is an important determinant of how an individual perceives the risks from COVID-19 as well as the policy responses to address the virus. This adds important insights for the mechanisms driving our findings in section II and, more broadly, the mechanisms that drive social network effects on behavior documented in previous works. Indeed, friend exposure shapes individuals' beliefs about COVID-19 and the need for public-health-motivated restrictions on public life, providing evidence for an important beliefs-based channel that in turn affects mobility behavior.

It is noteworthy that we find these effects in a setting in which publicly available information from domain experts was ubiquitous. This suggests that the effects on beliefs are not primarily the result of friends conveying information that is otherwise hard to access. It is instead more consistent with a mechanism whereby information resonates more with individuals when it is communicated by friends. For instance, Malmendier and Veldkamp (2022) propose a model of learning in which people process the same information differently depending on who delivers it. In this model, "abstractly learned statistics and other information tends to be weighted significantly less than information gathered from . . . the experiences of others whom we care about, identify with or empathize with."

IV. Evidence on Friend-Exposure Effects from Public Data

In this final section, we briefly describe analyses that confirm our main results using publicly available ZCTA-level data. Appendix B provides more details and the complete set of results.[27]

[26] In addition to the results presented in this section, in tables A.20, A.21, and A.22, we study heterogeneities in the observed effects of friend exposure to COVID-19, finding results largely consistent with those presented in sec. II.C.

[27] We provide replication code to reproduce the analyses in this section through a repository on Dataverse (Bailey et al. 2023). Some of the results depend on data from SafeGraph, which cannot be included in the replication package itself, but which can be easily accessed

For this analysis, we combine public data on mobility from SafeGraph with the Social Connectedness Index data from Facebook (see Bailey et al. 2018b). We find that social distancing in a ZCTA increases when COVID-19 exposure increases in other locations with many social links to the target ZCTA. While this analysis does not allow us to control for many individual-level characteristics that are correlated with changes in social distancing behavior and exposure to COVID-19, it has the advantage that the SafeGraph mobility data are based on a different and larger set of individuals, thus mitigating concerns that the results discussed in the main body are merely an artifact of the somewhat selected sample of Facebook users who have consented to sharing and storing their location information.

We also disaggregate the SafeGraph mobility data by point of interest and merchant type to understand which types of establishments are visited less often by individuals with high friend exposure to COVID-19. Using a difference-in-differences analysis similar to section II, we document that individuals living in places that are socially connected to highly exposed places tend to disproportionately reduce discretionary visits to places requiring social interaction with others. There are smaller and insignificant effects on less discretionary visits, such as those to food and beverage stores and health care providers.

V. Conclusion

We use de-identified data from Facebook to show that personal connections to COVID-19 hot spots significantly affected individuals' social distancing behavior during the COVID-19 pandemic. At the onset of the pandemic, individuals whose friends lived in areas with worse coronavirus outbreaks reduced their mobility more than their otherwise similar neighbors with fewer friends in affected areas. As the pandemic spread across the United States, users with more friends in emerging hot spots in one month continued to reduce their mobility in that month relative to their neighbors with friends in other parts of the country. Analyzing mobility at the individual level in such a changes-on-changes specification allows us to rule out various confounds when establishing the effect of friend experiences on social distancing behavior.

We then use data on public Facebook posts and group memberships to show that friend exposure to COVID-19 cases affects individuals' stated beliefs about the risks of COVID-19 and the benefits of mitigating public health behavior. Specifically, users with higher friend exposure to COVID-19 cases are more likely to post about the coronavirus and are less likely to

by researchers who agree to the data's terms of use. Instructions on downloading and organizing the data are available in the replication package.

oppose distancing in these posts. These users are also less likely to join Facebook groups advocating for a reopening of the economy.

A key conclusion of our work is that friend experiences affected beliefs about the COVID-19 pandemic at a time when information from many expert sources was ubiquitous, and when COVID-19 received unparalleled press coverage and public messaging. It is thus unlikely that the main reason why friend experiences were so influential is that they provided a low-cost source of information. Instead, it is more likely that information received from friends resonates particularly with people, and thus receives a substantial weight in the belief formation process. Our results therefore add new insight into the general mechanism underlying the important role of social networks in shaping individuals' beliefs and subsequent actions. We believe that studying both the empirical and theoretical properties of such an "information resonance" channel is a very promising area for future research. It is also important to highlight that under such mechanisms, friend experiences are likely to influence beliefs and behavior in both desirable and undesirable ways—consistent, for example, with the role of social interactions in spreading conspiracy theories—with a limited role of providing expert information as a countervailing force. This insight can play an important role in helping policymakers design more effective public information campaigns across a range of settings, from public health to consumer protection.

Data Availability

The analysis code used in the article can be found in Bailey et al. (2023) in the Harvard Dataverse, https://doi.org/10.7910/DVN/QYZLHT. In the instructions file, we describe how to download the data for the aggregated analyses presented in appendix B. The data used in the other analyses are proprietary and cannot be accessed without a research agreement with Meta.

References

Alekseev, Georgij, Safaa Amer, Manasa Gopal, Theresa Kuchler, J. W. Schneider, Johannes Stroebel, and Nils Wernerfelt. 2022. "The Effects of COVID-19 on U.S. Small Businesses: Evidence from Owners, Managers, and Employees." *Management Sci.* 69 (1): 7–24.

Allcott, Hunt, Levi Boxell, Jacob Conway, Billy A. Ferguson, Matthew Gentzkow, and Benny Goldman. 2020a. "What Explains Temporal and Geographic Variation in the Early US Coronavirus Pandemic?" Working Paper no. 27965 (October), NBER, Cambridge, MA.

Allcott, Hunt, Levi Boxell, Jacob Conway, Matthew Gentzkow, Michael Thaler, and David Yang. 2020b. "Polarization and Public Health: Partisan Differences in Social Distancing during the Coronavirus Pandemic." *J. Public Econ.* 191:104254.

Armantier, Olivier, Wändi Bruine de Bruin, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar. 2015. "Inflation Expectations and Behavior: Do Survey Respondents Act on Their Beliefs?" *Internat. Econ. Rev.* 56 (2): 505–36.

Armona, Luis, Andreas Fuster, and Basit Zafar. 2019. "Home Price Expectations and Behaviour: Evidence from a Randomized Information Experiment." *Rev. Econ. Studies* 86 (4): 1371–410.

Bachmann, Rüdiger, Tim O. Berg, and Eric R. Sims. 2015. "Inflation Expectations and Readiness to Spend: Cross-Sectional Evidence." *American Econ. J. Econ. Policy* 7 (1): 1–35.

Bailey, Michael, Rachel Cao, Theresa Kuchler, and Johannes Stroebel. 2018a. "The Economic Effects of Social Networks: Evidence from the Housing Market." *J.P.E.* 126 (6): 2224–76.

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

Bailey, Michael, Eduardo Dávila, Theresa Kuchler, and Johannes Stroebel. 2019. "House Price Beliefs and Mortgage Leverage Choice." *Rev. Econ. Studies* 86 (6): 2403–52.

Bailey, Michael, Patrick Farrell, Theresa Kuchler, and Johannes Stroebel. 2020a. "Social Connectedness in Urban Areas." *J. Urban Econ.* 118:103264.

Bailey, Michael, Abhinav Gupta, Sebastian Hillenbrand, Theresa Kuchler, Robert Richmond, and Johannes Stroebel. 2021. "International Trade and Social Connectedness." *J. Internat. Econ.* 129:103418.

Bailey, Michael, Drew Johnston, Martin Koenen, Theresa Kuchler, Dominic Russel, and Johannes Stroebel. 2022a. "The Social Integration of International Migrants: Evidence from the Networks of Syrians in Germany." Working Paper no. 29925 (April), NBER, Cambridge, MA.

———. 2023. "Replication Data for: 'Social Networks Shape Beliefs and Behaviors: Evidence from Social Distancing during the COVID-19 Pandemic.'" Harvard Dataverse, https://doi.org/10.7910/DVN/QYZLHT.

Bailey, Michael, Drew Johnston, Theresa Kuchler, Dominic Russel, Johannes Stroebel, and Bogdan State. 2020b. "The Determinants of Social Connectedness in Europe." In *Social Informatics: 12th International Conference*, edited by Samin Aref. Cham, Switzerland: Springer.

Bailey, Michael, Drew Johnston, Theresa Kuchler, Johannes Stroebel, and Arlene Wong. 2022b. "Peer Effects in Product Adoption." *American Econ. J. Appl. Econ.* 14 (3): 488–526.

Bakkensen, Laura A., and Lint Barrage. 2022. "Going Underwater? Flood Risk Belief Heterogeneity and Coastal Home Price Dynamics." *Rev. Financial Studies* 35 (8): 3666–709.

Banerjee, Abhijit, Arun G. Chandrasekhar, Esther Duflo, and Matthew O. Jackson. 2019. "Using Gossips to Spread Information: Theory and Evidence from Two Randomized Controlled Trials." *Rev. Econ. Studies* 86 (6): 2453–90.

Bargain, Olivier, and Ulugbek Aminjonov. 2020. "Trust and Compliance to Public Health Policies in Times of COVID-19." *J. Public Econ.* 192:104316.

Barrios, John M., Efraim Benmelech, Yael V. Hochberg, Paola Sapienza, and Luigi Zingales. 2021. "Civic Capital and Social Distancing during the Covid-19 Pandemic." *J. Public Econ.* 193:104310.

Barrios, John M., and Yael V. Hochberg. 2021. "Risk Perceptions and Politics: Evidence from the COVID-19 Pandemic." *J. Financial Econ.* 142 (2): 862–79.

Bordalo, Pedro, Giovanni Burro, Katherine B. Coffman, Nicola Gennaioli, and Andrei Shleifer. 2022. "Imagining the Future: Memory, Simulation and Beliefs about COVID." Working Paper no. 30353 (August), NBER, Cambridge, MA.

Brodeur, Abel, David Gray, Anik Islam, and Suraiya Bhuiyan. 2021. "A Literature Review of the Economics of COVID-19." *J. Econ. Surveys* 35 (4): 1007–44.

Brodeur, Abel, Idaliya Grigoryeva, and Lamis Kattan. 2021. "Stay-at-Home Orders, Social Distancing, and Trust." *J. Population Econ.* 34 (4): 1321–54.

Brzezinski, Adam, Valentin Kecht, David Van Dijcke, and Austin L Wright. 2021. "Science Skepticism Reduced Compliance with COVID-19 Shelter-in-Place Policies in the United States." *Nature Human Behaviour* 5 (11): 1519–27.

Bursztyn, Leonardo, Aakaash Rao, Christopher Roth, and David Yanagizawa-Drott. 2023. "Opinions as Facts." *Rev. Econ. Studies* 90 (4): 1832–64.

Charoenwong, Ben, Alan Kwan, and Vesa Pursiainen. 2020. "Social Connections with COVID-19–Affected Areas Increase Compliance with Mobility Restrictions." *Sci. Advances* 6 (47): eabc3054.

Chetty, Raj, Matthew O. Jackson, Theresa Kuchler, Johannes Stroebel, Nathaniel Hendren, Robert B. Fluegge, Sara Gong, et al. 2022a. "Social Capital I: Measurement and Associations with Economic Mobility." *Nature* 608 (7921): 108–21.

———. 2022b. "Social Capital II: Determinants of Economic Connectedness." *Nature* 608 (7921): 122–34.

Christakis, Nicholas A., and James H. Fowler. 2007. "The Spread of Obesity in a Large Social Network over 32 Years." *New England J. Medicine* 357 (4): 370–79.

———. 2008. "The Collective Dynamics of Smoking in a Large Social Network." *New England J. Medicine* 358 (21): 2249–58.

D'Acunto, Francesco, Daniel Hoang, Maritta Paloviita, and Michael Weber. 2023. "IQ, Expectations, and Choice." *Rev. Econ. Studies* 90 (5): 2292–325.

Dingel, Jonathan I., and Brent Neiman. 2020. "How Many Jobs Can Be Done at Home?" *J. Public Econ.* 189:104235.

Dong, Ensheng, Hongru Du, and Lauren Gardner. 2020. "An Interactive Web-Based Dashboard to Track COVID-19 in Real Time." *Lancet Infectious Diseases* 20 (5): 533–34.

Facebook. 2020. "2019 Annual Report on Form 10-K." Facebook, Menlo Park, CA.

Fletcher, Jason M., and Stephen L. Ross. 2018. "Estimating the Effects of Friends on Health Behaviors of Adolescents." *Health Econ.* 27 (10): 1450–83.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus. 2021a. "Five Facts about Beliefs and Portfolios." *A.E.R.* 111 (5): 1481–522.

———. 2021b. "The Joint Dynamics of Investor Beliefs and Trading during the COVID-19 Crash." *Proc. Nat. Acad. Sci. USA* 118 (4): e2010316118.

Goodman-Bacon, Andrew. 2021. "Difference-in-Differences with Variation in Treatment Timing." *J. Econometrics* 225 (2): 254–77.

Greenwood, Shannon, Andrew Perrin, and Maeve Duggan. 2016. "Social Media Update 2016." Report, Pew Res. Center, Washington, DC.

Guiso, Luigi, Giorgio Gulino, and Ruben Durante. 2020. "Civic Capital and Social Distancing: Evidence from Italians' Response to COVID-19." https://cepr.org/voxeu/columns/civic-capital-and-social-distancing-evidence-italians-response-covid-19.

Holtz, David, Michael Zhao, Seth G. Benzell, Cathy Y. Cao, Mohammad Amin Rahimian, Jeremy Yang, Jennifer Allen, et al. 2020. "Interdependence and the Cost of Uncoordinated Responses to COVID-19." *Proc. Nat. Acad. Sci. USA* 117 (33): 19837–43.

Hu, Zhongchen. 2022. "Social Interactions and Households' Flood Insurance Decisions." *J. Financial Econ.* 144 (2): 414–32.

Huang, Grace C., Jennifer B. Unger, Daniel Soto, Kayo Fujimoto, Mary Ann Pentz, Maryalice Jordan-Marsh, and Thomas W. Valente. 2014. "Peer Influences:

The Impact of Online and Offline Friendship Networks on Adolescent Smoking and Alcohol Use." *J. Adolescent Health* 54 (5): 508–14.
Hutto, Clayton J., and Eric Gilbert. 2014. "VADER: A Parsimonious Rule-Based Model for Sentiment Analysis of Social Media Text." *Proc. Internat. AAAI Conference Web and Social Media* 8 (1): 216–25.
Iyer, Shankar, Brian Karrer, Daniel Citron, Farshad Kooti, Paige Maas, Zeyu Wang, Eugenia Giraudy, Ahmed Medhat, P. Alex Dow, and Alex Pompe. 2023. "Large-Scale Measurement of Aggregate Human Colocation Patterns for Epidemiological Modeling." *Epidemics* 42 (March): 100663.
Jacobson, Robert M., Jennifer L. St. Sauver, and Lila J. Finney Rutten. 2015. "Vaccine Hesitancy." *Mayo Clinic Proc.* 90 (11): 1562–68.
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.
Kuchler, Theresa, Yan Li, Lin Peng, Johannes Stroebel, and Dexin Zhou. 2022. "Social Proximity to Capital: Implications for Investors and Firms." *Rev. Financial Studies* 35 (6): 2743–89.
Kuchler, Theresa, Dominic Russel, and Johannes Stroebel. 2022. "JUE Insight: The Geographic Spread of COVID-19 Correlates with the Structure of Social Networks as Measured by Facebook." *J. Urban Econ.* 127 (January): 103314.
Kuchler, Theresa, and Basit Zafar. 2019. "Personal Experiences and Expectations about Aggregate Outcomes." *J. Finance* 74 (5): 2491–542.
Maas, Paige, Shankar Iyer, Andreas Gros, Wonhee Park, Laura McGorman, Chaya Nayak, and P. Alex Dow. 2019. "Facebook Disaster Maps: Aggregate Insights for Crisis Response and Recovery." In *Proceedings of the 16th International Conference on Information Systems for Crisis Response and Management, Valencia (Spain), May 19–22*, edited by Zeno E. Franco, José J. González, and José H. Canós, 836–47. Valencia: Univ. Politécnica Valencia.
Makridis, Christos, and Tao Wang. 2020. "Learning from Friends in a Pandemic: Social Networks and the Macroeconomic Response of Consumption." https://papers.ssrn.com/abstract=3601500.
Malmendier, Ulrike, and Stefan Nagel. 2011. "Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?" *Q.J.E.* 126 (1): 373–416.
Malmendier, Ulrike, and Laura Veldkamp. 2022. "Information Resonance." Working paper, Econometrics Laboratory, Univ. California, Berkeley.
Mayer, Maximilian. 2023. "Climate Change Concerns and Information Spillovers from Socially-Connected Friends." IWH Discussion Paper no. 2/2023, Leibniz Inst. Wirtschaftsforschung Halle.
Peretti-Watel, Patrick, Valérie Seror, Sébastien Cortaredona, Odile Launay, Jocelyn Raude, Pierrea Verger, Lisa Fressard, et al. 2020. "A Future Vaccination Campaign against COVID-19 at Risk of Vaccine Hesitancy and Politicisation." *Lancet Infectious Diseases* 20 (7): 769–70.
Pew Research Center. 2020. "U.S. Public Sees Multiple Threats from the Coronavirus—and Concerns Are Growing." Report, Pew Res. Center, Washington, DC.
Rasul, Imran, and Paola Giuliano. 2020. "Compliance with Social Distancing during the Covid-19 Crisis." https://cepr.org/voxeu/columns/compliance-social-distancing-during-covid-19-crisis.
Ratnadiwakara, Dimuthu. 2021. "Flooded Social Connections." *Q. J. Finance* 11 (4): 2150018.
Rehbein, Oliver, and Simon Rother. 2020. "Distance in Bank Lending: The Role of Social Networks." ECONtribute Discussion Paper no. 033, Reinhard Selten Inst., Bonn and Cologne.

Roth, Christopher, and Johannes Wohlfart. 2020. "How Do Expectations about the Macroeconomy Affect Personal Expectations and Behavior?" *Rev. Econ. and Statis.* 102 (4): 731–48.

Rothwell, Jonathan T., Christos Makridis, Christina Ramirez, and Sonal Desai. 2021. "COVID-19 Information Treatments Shift Related Policy Preferences and Plans for Consumer Behavior." https://papers.ssrn.com/abstract=3859751.

Sato, Ryoko, and Yoshito Takasaki. 2019. "Peer Effects on Vaccination Behavior: Experimental Evidence from Rural Nigeria." *Econ. Development and Cultural Change* 68 (1): 93–129.

Schwartz, Joe. 2018. "Bing Maps Tile System." Report, Microsoft, Redmond, WA.

Simonov, Andrey, Szymon Sacher, Jean-Pierre Dubé, and Shirsho Biswas. 2022. "Frontiers: The Persuasive Effect of Fox News: Noncompliance with Social Distancing during the COVID-19 Pandemic." *Marketing Sci.* 41 (2): 230–42.

Tian, Yuan, Maria Esther Caballero, and Brian K. Kovak. 2022. "Social Learning along International Migrant Networks." *J. Econ. Behavior and Org.* 195:103–21.

Weber, Elke U., and Paul C. Stern. 2011. "Public Understanding of Climate Change in the United States." *American Psychologist* 66 (4): 315–28.

Wilson, Riley. 2022. "The Impact of Social Networks on EITC Claiming Behavior." *Rev. Econ. and Statis.* 104 (5): 929–45.

Xu, Yilan, and Sebastien Box-Couillard. 2023. "Social Learning about Climate Change Risk." Working paper, Dept. Agricultural and Consumer Econ., Univ. Illinois, Urbana-Champaign.