

# Friends Don't Let Friends Free Ride

Nicholas Eubank\* and Dorothy Kronick†

January 5, 2021

## Abstract

Theory predicts that social sanctioning can solve the collective action problem, but only when people find out whether their peers participate. We evaluate this prediction using data from the near-universe of cell-phone subscribers in Venezuela. Those whose behavior is more easily observed by peers are much more likely to protest and much more likely to sign a political petition than otherwise similar people in less-visible social network positions. Together with qualitative and survey data, we interpret this finding as evidence that social network structure can facilitate (or frustrate) social sanctioning as a solution to the collective action problem.

Word count: 9,844

---

\*Duke Social Science Research Institute, [nick@nickeubank.com](mailto:nick@nickeubank.com)

†Department of Political Science, University of Pennsylvania, [kronick@sas.upenn.edu](mailto:kronick@sas.upenn.edu)

‡We are grateful to Seth Bromberger and to contributors to open source projects like Julia, LightGraphs, iGraph, and Python. For comments, we thank Guy Grossman, Yue Hou, Horacio Larreguy, Jenn Larson, Michele Margolis, John Marshall, Cyrus Samii, Zachary Steinert-Threlkeld, Dawn Teele, Stephanie Zonszein, Liz Zechmeister, and participants at the NYU Methods Workshop and NEWEPS. We also thank Vanderbilt University for generous access to server time. **Disclosure:** This work was made possible by a data-sharing agreement between the authors and a telecommunications firm. In exchange for access to the data, the authors provided training in coding and data analysis. At no time did the authors provide the company with any financial compensation, nor did the company provide the authors with any financial compensation.

Fifty-five years ago, Mancur Olson (1965) posed an enduring question. Why do people protest, or lobby, or strike, or otherwise take part in collective action, given that doing so requires effort and risk, while everyone—participant and free-rider alike—reaps the spoils of victory?

One influential explanation is that free riding is rarely free: those who sit out often suffer social sanction, while participants enjoy social rewards (Ostrom, 1990). Indeed, a large experimental literature reveals that social influence affects whether people engage in collective political activities like signing petitions (Paler et al., 2018), attending protests (Enikopolov et al., 2017b), and voting (Nickerson, 2008; Gerber et al., 2008, 2016; Panagopoulos, 2010).

Yet we know little about why certain communities are better able to wield this social influence, or why certain individuals are more susceptible to it.

Theory points to an important role for *information capital*: the ability to acquire and/or spread information (Jackson, 2019, 318). The idea is simple. Social influence can only solve the collective action problem if people find out whether others participate, which requires a social network conducive to information diffusion. This prediction—that information diffusion enables community enforcement—is central to influential theories of prosocial behavior (Larson, 2017b; Wolitzky, 2012; Jackson et al., 2017, 75).

If true, this prediction would provide considerable insight into the puzzle of collective political action. We would better understand why some communities mobilize easily while others flounder. But empirical tests have been limited. For one thing, social network structure is difficult to observe except in small villages (Larson and Lewis, 2017, 2020; Eubank et al., 2018; Ferrali et al., 2019) or online (Larson, 2017a; Steinert-Threlkeld, 2017). For another, *communication centrality*—Jackson’s preferred measure of information capital—is computationally difficult to estimate (2019, 320). Perhaps most challenging of all is observing the outcome: who participates in collective political action.

We address these challenges with original data from Venezuela. The data allow us to map the offline social network, estimate each person’s communication centrality, and observe who participates in two forms of collective action: a petition, and a related protest. Using these data, we find that people with higher communication centrality are much more likely to sign the petition and much more likely to protest than otherwise similar people in less-visible social network positions. We then use features of the Venezuelan context to argue that the mechanism linking communication centrality to political participation is *social influence*, i.e., peer sanctioning and reward. We conclude that people with high communication centrality are more susceptible to social influence, which enables their communities to enforce participation. Information capital distinguishes participants from free-riders.

We map the structure of Venezuela’s social network using data from the near-

universe of cell-phone subscribers.<sup>1</sup> To measure *communication centrality*, we introduce a simulation-based approximation; simple statistics such as number of immediate connections do not capture the speed of information diffusion (Newman et al., 2006; Banerjee et al., 2013).<sup>2</sup> We then observe two forms of political participation: whether each person signs a petition demanding a referendum on recalling Venezuelan President Nicolás Maduro (based on a list of signatories), and whether each person participates in a protest aimed at convincing the government to honor the petition (based on cell phone location during the protest). These measures are behavioral, not self-reported.

We find that communication centrality is correlated with political participation: people with high levels of communication centrality are much more likely to participate—in the protest and in petition signing—than people with lower levels of communication centrality. This is true even when we compare participants to non-participants who vote at the same polling place, have the same party registration, are of the same gender and similar age, and have a similar level of geographic mobility. Within this matched sample, a one-standard-deviation increase in communication centrality predicts a 2.8-percentage-point increase in the probability of signing the petition. Placebo tests suggest that these relationships are not driven by unobserved characteristics that differ across participants and non-participants. Our analysis is descriptive, but we do compare participants with non-participants who are more observationally similar than the comparison sets used in previous descriptive work (Larson et al., 2019; Steinert-Threlkeld, 2017).

We then exploit features of the Venezuelan context—such as the fact that protests were organized publicly, rather than clandestinely—to argue that the mechanism linking communication centrality to political participation is *social sanctioning*, rather than (e.g.) awareness (Christensen and Garfias, 2018) or mirroring (Siegel, 2009; Rolfe, 2012).<sup>3</sup> For example, using responses to original survey questions, we find that more than one third of respondents had friends or family members who weighed in on others’ participation decisions.

Our analysis of petition signing also reveals a risk of using protests in the study of social networks and political participation. The relationship between communication centrality and political participation is much stronger for protest attendance than for petition signing. Based in part on responses to original survey questions, we attribute this result to the fact that a protest is an inherently social activity, usually attended in the company of friends; thus,

---

<sup>1</sup>See Eagle et al. (2009) on inferring social ties from cell-phone meta-data.

<sup>2</sup>Our approximation can be applied to any network data; we contributed the code to the open-source `LightGraphs.jl` library, and it is available [here](#). Our own full implementation will be posted upon publication of this paper, and is available by request in the meantime.

<sup>3</sup>We do not directly test models like Rolfe (2012) or Siegel (2009) because (to the best of our knowledge) doing so requires comparison across partitioned networks (e.g. Eubank et al., 2018), whereas we observe a single large network.

protesters might have higher communication centrality than non-protesters simply because they are more sociable. In other words, they might also be more likely to attend (say) dinner parties. This finding sounds a note of caution for related empirical work (c.f. Gonzalez, 2017; Larson et al., 2019): network position and protest participation might be correlated for reasons that have nothing to do with solving collective action problems.

In addition to the networks literature cited above, our results contribute to literature on the role of social networks in Latin American politics. It has long been understood that social ties play an important role in mediating access to excludable benefits (Calvo and Murillo, 2013; Stokes et al., 2013), screening prospective clients (Auyero, 2000; Ujhelyi and Calvo, 2011), and voter coordination (Arias et al., 2019); we build on these insights in studying the relationship between social networks and protest. Our focus on network structure complements recent work highlighting “how much is at stake” in a given protest or election (Aytaç and Stokes, 2019, 38–39): social networks that enable peer sanctioning will amplify the effects of event characteristics that spark widespread participation.

Overall, we provide (to the best of our knowledge) the first quantitative empirical evidence in favor of theories such as Larson (2017b) and Jackson et al. (2017): social networks conducive to information diffusion can help solve the collective action problem. Even compared to others in the same small neighborhood, of the same political persuasion, and of the same age and gender, Venezuelans more exposed to peer pressure were less likely to free ride on their compatriots’ costly actions against an increasingly authoritarian regime.

## 1 Context: Collective Action in Venezuela

In the two years after Nicolás Maduro took office as president of Venezuela in April of 2013, the country suffered the worst recession in its recorded history (Kronick, 2015).<sup>4</sup> Maduro’s approval ratings dipped below 25% (Datanalisis, 2016), and, in April of 2016, a coalition of opposition parties decided to collect signatures in support of holding a recall referendum: a national yes-or-no vote on whether to recall Maduro from office. Under the Venezuelan constitution, written during the presidency of Maduro’s predecessor and mentor Hugo Chávez, the signatures of 1% of registered voters would suffice to begin the recall referendum process.<sup>5</sup>

The coalition of opposition parties set up petition-signing stations throughout

---

<sup>4</sup>Since then, the recession has become the worst in recorded Latin American history.

<sup>5</sup>Technically, it is an electoral regulation (not the Constitution) that requires the signatures of 1% of registered voters to begin the process. The regulation specifies that, with signatures of 1% of the electorate in hand, the Electoral Council would supervise the collection of signatures of the 20% of the electorate required by the Constitution in order to hold the recall vote.

the country, where registered voters could add their signatures and thumbprints to the petition (see *Noticia al Día*, 2016, for photos and a description). The stations were open for two days (April 27 and April 28, 2016); on May 2, opposition leaders announced that they had collected and submitted to the National Electoral Council 1.8 million signatures—nearly ten times the required 1% (Toro, 2016).

The Electoral Council claimed that the signatures were not valid. In the view of economist Francisco Rodríguez, “There was just nothing even resembling a normally coherent argument about why it was that the referendum was stopped. The government alleges that there was fraud in the collection of signatures, but . . . there were enough signatures, even excluding the presumed fraudulent signatures, to get the process to go forward. But nevertheless the government stopped it” (2017).

Opposition leaders responded by calling for a large demonstration, the *Toma de Caracas* (“taking of Caracas”), and they scheduled it for September 1. The time and location of the march were publicized on Twitter and in major media outlets, and the march itself was covered in the national and international press. The protest organizers claimed that more than one million people participated in the march; an independent measure based on photographs estimated participation at 700,000 (Rodríguez, 2016).

We use (a) signing of the recall referendum petition and (b) presence at the *Toma de Caracas* as our measures of participation in collective political action. (See Section 3 for measurement details.) Both were costly activities subject to the collective action problem. In addition to the time and transportation costs required to sign the petition, signatories had reason to fear retribution: signatories of an earlier, similar petition had been subsequently discriminated against in the labor market (Hsieh et al., 2011) and in applications for government benefits (Albertus, 2015). Likewise, in addition to the time and transportation costs required to attend the protest, participants had reason to fear violence: though not common, participants in earlier protests had suffered injuries, arrest, torture, and even fatalities (Toro, 2014).

## 2 Theory: Networks and Political Participation

The free-rider problem plagues protests, petitions, and other collective political activities: why spend time (and risk reprisals) participating when you could instead stay home and still benefit from any political gains?

One answer is social sanctioning. A community agreement to shun, chastise, or otherwise punish shirkers provides powerful incentives against free-riding (Ostrom, 1990; Kandori, 1992; Fearon and Laitin, 1996). Likewise, social

recognition of people who do “do their civic duty” increases the returns to participation (Sinclair, 2012).

Jackson et al. (2017, 75) explain why information diffusion underpins any system of community enforcement:

If friends and neighbors are made quickly aware of the individual’s behavior, then they can react quickly. This is important for providing the right incentives, as the threat of a punishment in the near future will have the greatest scope for disciplining behavior. If instead, the setting is such that it takes a long time for one’s friends to learn of misbehavior then it becomes difficult to provide incentives for individuals to behave according to some desired social norm.

Theorists have formalized this intuition. Larson (2017b), for example, presents a model in which the threat of collective sanction of defectors sustains cooperation in one-shot bilateral prisoner’s dilemma games (played between random pairs of individuals) (à la Kandori, 1992). In the equilibrium of interest, cooperation can be sustained only among individuals whose network positions ensure that, if they defect, word will spread quickly enough to ensure a credible threat of sanction by future partners. This excludes people in isolated network positions, for whom the threat of collective sanction is weak.<sup>6</sup>

These theories yield clear comparative statics: people in more visible network positions—that is, network positions that facilitate the spread of information about them—are more likely to participate, because they face stronger potential social sanction for staying home.

Of course, information diffusion might enable collective action through means other than social sanctioning. Most obviously, it might increase awareness of a protest or other political activity, or it might facilitate coordination among protest organizers evading government disruption (Little, 2015; Christensen and Garfias, 2018; Enikopolov et al., 2017a). In our context, however, organizers were not forced to promote events solely via private social networks; on the contrary, the petition signature drive and the *Toma de Caracas* protest were announced on social media and covered extensively in major press outlets. Private coordination through social networks therefore did not play the essential logistical role that it played in, for example, Arab Spring protests (Steinert-Threlkeld, 2017; Steinert-Threlkeld et al., 2015). The very public nature of these activities motivates our focus on social sanctioning (or social approbation) as the mechanism of interest.

Similarly, the presence of (at the time) reliable public opinion polls motivates

---

<sup>6</sup>A similar dynamic emerges in a repeated public goods game in Wolitzky (2012).

our choice to focus on free riding rather than preference falsification as the impediment to collective action. Preference falsification can generate the distorted impression of widespread support for the regime, which then hinders protest participation (Kuran, 1991; Steinert-Threlkeld, 2018; Little, 2015).<sup>7</sup> But in Venezuela, media outlets regularly reported the president’s unpopularity: in the summer of 2016, just before the *Toma de Caracas* protest, major pollsters placed Maduro’s approval ratings between 22% and 31%. There was no public perception of widespread support for the government.

We do not seek to study the direct influence of one person’s observed or expected participation on her friends’ participation decisions (as in, e.g., Rolfe, 2012; Siegel, 2009; Banerjee et al., 2014; Eubank et al., 2018; Ferrali et al., 2019; Cantoni et al., 2019). Rather, our interest lies in how social network structure “determines the forms of cooperative behavior that can be maintained” (Jackson et al., 2017, 5) by shaping the *anticipated* social cost of non-participation—which exists even (or especially) when sanctioning never actually occurs in equilibrium. Our empirical analysis is thus more in line with Jackson et al. (2012, 1882–1883) or Larson et al. (2019), both of which compare the network positions of cooperative and non-cooperative agents.

### 3 Data: Communication Centrality and Participation

The theories summarized in Section 2 predict that, all else equal, people with higher communication centrality are more susceptible to peer pressure and therefore more likely to participate in collective action. To study this prediction empirically, we implement a measure of communication centrality—or, in other words, a measure of how quickly information about a person diffuses through her social network.

#### 3.1 Mapping the Social Network

We first map the social network through which information diffuses, using cell-phone meta-data from a Venezuelan telecommunications company (“Partner

---

<sup>7</sup>If we were to focus on preference falsification, we would have used a different network model to motivate our analysis, such as complex contagion (Centola, 2018; Centola and Macy, 2007) or social context (Siegel, 2009). In these models, participation snowballs: individuals choose to protest as they observe peers protesting, because these observations lead them to update about the level of support for the government in the population. Our focus instead on free riding motivates our choice of Larson (2017b) and Wolitzky (2012) as theoretical foundations for the analysis.

Telecom”).<sup>8</sup>

The data include records of all SMS (text) and voice transactions sent or received by subscribers of the Partner Telecom from June 2016 to February 2017, totaling approximately 30 billion transactions. Because we observe transactions sent *or* received by Partner Telecom subscribers, our data include information about other providers’ clients, too. Each record includes type (voice or text), identifiers for both caller and receiver, date and time, duration, and GPS coordinates of the antenna tower through which all calls were placed (we do not observe antenna towers for text messages).<sup>9</sup>

Within this data, we classify two users as *connected* if they (a) call each other at least twice *or* text each other at least twelve times<sup>10</sup> in (b) at least two of the eight months in our data. Results are similar using call thresholds of four voice calls or twenty-four texts and six voice calls or thirty-six texts. In our baseline specification, our network has about 27 million individual subscribers (vertices in the network), the median person has nine immediate connections, and the average person has 9.3 immediate connections.

In principle, we could use call frequency to measure the strength of ties (rather than the mere presence of a tie). In practice, we would not trust this measure. First, it is unclear whether call frequency is strongly correlated with the social importance of a relationship. Second, phone calls are only one of several modes of communication. Even if overall communication frequency were indicative of social importance, weighting by *call* frequency may underestimate the importance of relationships among individuals who live nearby and primarily communicate face-to-face—or those who communicate via email or WhatsApp.<sup>11</sup> See Appendix A for additional discussion.

Our mapping of the social network is imperfect. First, while we do observe communication between Partner Telecom subscribers and those using other providers, we do not observe connections among pairs of users when neither is a Partner Telecom subscriber. (To preserve anonymity, we cannot report their market share.) Second, our data may generate false positives: a person might make many calls to a local business, for example, without being socially connected to that business (though we do filter out one-off business calls by requiring connections in at least two of our eight months of data, and we

---

<sup>8</sup>The telecommunications company requested anonymity.

<sup>9</sup>All identified data is stored on a computer with no physical means of connecting to the internet (an air-gapped computer), in an access-controlled room, on encrypted hard drives; the data protection protocols were approved by two Institutional Review Boards.

<sup>10</sup>This reflects the fact that messages are about six times more common than voice calls.

<sup>11</sup>WhatsApp looks like data usage to telecommunications firms, so our data excludes it.

filter out very large businesses by dropping the 1% of users with the most connections). Likewise, we cannot avoid false negatives even among Partner Telecom subscribers; family members or others who see each other daily may rarely call or text.

In our view, the advantages of our data outweigh these drawbacks. Many studies map social networks by asking people for the names of their closest friends (e.g. Larson and Lewis, 2017; Banerjee et al., 2014; Rojo and Wibbels, 2014; Fafchamps and Vicente, 2013; Dionne, 2015). Our data provide more breadth: in our primary specification, we observe 27 million unique phone lines, or the near universe of lines in the country.<sup>12</sup> Our data also provide more depth, capturing even weak social ties that may be censored when people are asked to list a finite number (often five) of their closest friends—weak ties that have been found to be important (Granovetter, 1973). Third, our data measure actual, non-public, offline communication patterns. They therefore reveal connections that might be under-reported in surveys or absent online, like inter-class or inter-party ties.

These advantages have led economists and computer scientists to use cell-phone meta-data for studying phenomena like migration and mobility (Blumenstock, 2012; Wesolowski et al., 2012), socio-economic status (Blumenstock et al., 2015), the geography of network communities (Blondel et al., 2010; Barthélemy, 2011), and knowledge diffusion (Björkegren, 2018). To the best of our knowledge, we are the first to use cell-phone meta-data to study political participation.

## 3.2 Measuring Communication Centrality

With this mapping of the social network in place, we measure a person’s communication centrality by simulating an information diffusion process. Simulation is required because of the intractability of analytically accounting for “all the possible paths that information might take, and some end up overlapping, producing correlation in the chance that information makes it from one node to another” (Jackson, 2019, 8).

For each individual  $i$ , we simulate a diffusion process starting at  $i$  and record the number of  $i$ ’s peers who are reached (that is, informed) by the diffusion process at each time  $t$ .<sup>13</sup> The simulation proceeds as follows:

---

<sup>12</sup>The Venezuelan population was approximately 31 million in 2016; the World Bank estimates 93 cell lines per 100 people.

<sup>13</sup>In some theoretical models, information about a person *not* participating would have to spread without her help—i.e., in the network  $g \setminus i$ —while she might help spread information if she *did* participate. We abstract away from this distinction, allowing the information

1. At time  $t = 0$ , the subject vertex  $v_i$  is endowed with a unique piece of knowledge. In other words, she is *informed*. All other vertices are ignorant. Letting  $\mathbb{I}_{i,t}$  denote the set of vertices who have been informed by time  $t$  in a diffusion process beginning at  $v_i$ , this implies that  $\mathbb{I}_{i,0} = \{v_i\}$ .
2. At time  $t = 1$ , information spreads to each neighbor of  $v_i$  with probability  $p/\log(\text{degree}(v_i) + 1) \in (0, 1)$ . Our decision to normalize the probability of diffusion by the log of  $i$ 's number of neighbors—that is, by  $\log(\text{degree}(v_i) + 1)$ —is motivated by an empirical regularity: people with more friends do make more calls overall than users with fewer friends, but this increase is sub-linear in the number of friends (Miritello et al., 2013), likely due to time constraints (see also Larson and Lewis, 2017).<sup>14</sup>
3. From each vertex  $v_j$  informed at  $t = 1$  (that is, from each  $v_j \in \mathbb{I}_{i,1}$ ), the information spreads to  $v_j$ 's neighbors with probability  $p/\log(\text{degree}(v_j) + 1)$  in  $t = 2$ , creating a new set of informed vertices  $\mathbb{I}_{i,2}$ .
4. For each subsequent time period  $t$ , information spreads from each informed vertex  $v_j \in \mathbb{I}_{i,t-1}$  to each of  $v_j$ 's neighbors with probability  $p/\log(\text{degree}(v_j) + 1)$ , creating a new set of informed vertices  $\mathbb{I}_{i,t}$ .

We run this simulation 1,000 times for each person (vertex) in our samples of interest (described below). Averaging over the 1,000 simulations, we create two measures:

$$\begin{aligned} \text{Communication centrality: } N_{i,t}^g &\equiv |\mathbb{I}_{i,t}| \\ \text{Exposure to participants: } N_{i,t}^p &\equiv |\mathbb{I}_{i,t} \cap \mathbb{P}|, \\ &\text{where } \mathbb{P} \text{ is the set of participants} \end{aligned}$$

In other words, the *communication centrality* of person  $i$  at time  $t$  ( $N_{i,t}^g$ ) is the number of people informed about her behavior up to and including time  $t$ ; her *exposure to participants* is the number of other *participants* (people who signed the petition or attended the protest) informed about  $i$ 's behavior up to and including time  $t$ .

A key feature of communication centrality is that it captures the number of participants *and* non-participants reached by the diffusion process. If we only observed information diffusion to other participants, our results would be much more difficult to interpret. As we discuss in more detail below, any correlation between participation and *exposure to participants* could simply reflect political

---

diffusion process to begin at person  $i$  whether or not she participated. In our view, it is plausible that a person's peers could learn about her non-participation directly from her (for example, by asking, or indirectly through general awareness of a friend's schedule).

<sup>14</sup>Results are similar when we normalize by  $\text{degree}(v)$  rather than  $\log(\text{degree}(v) + 1)$ .

homophily: a tendency to associate with others who share a preference for (e.g.) attending protests. Communication centrality is instead a pure measure of network structure that does not consider whether or not the people reached by the diffusion process are political participants. Without it, we would not be able to rule out the possibility that what looks like social influence is in fact a shared preference for political participation. On the other hand, in light of theories suggesting that only participants have the moral authority to reward fellow participants or sanction free-riders, we also consider *exposure to participants* below.

No simple statistic—such as average degree or average shortest path length—captures the speed of information diffusion through a network (Newman et al., 2006). Moreover, the relationship between information diffusion and proxies for network structure in the literature—such as membership in civic organizations or self-reported number of friends—is generally unknowable.

Appendix Table C.3 reveals that communication centrality is highly correlated with but not identical to *eigenvector centrality*, a common measure that calculates the centrality of a node as the (scaled) sum of the centrality of its neighbors (Jackson, 2008, 41). To the best of our knowledge, this correlation has not been estimated before. If the correlation we observe ( $\approx 0.7$ ) were replicated in other data sets, analysis of eigenvector centrality would take on new social and political interpretations. Conversely, a low correlation in other network data sets would underscore the need for careful consideration of when to use communication centrality (which is more computationally demanding) and when to use eigenvector centrality (among other possible choices).<sup>15</sup>

### 3.3 Measuring Political Participation

Measuring petition signing is straightforward. The (identified) list of 1.7 million signatories was circulated; we merge the list (on national I.D. number) with our measures of communication centrality.

To measure protest activity, we classify as a *protester* any cell-phone subscriber with at least one call routed through a cell tower along the protest route during the time of the protest—except those who live or work nearby. Specifically, we exclude from the sample those who live or work in parishes adjacent to the protest route (parish is a sub-county administrative unit). We do this in order to reduce the incidence of false positives: the imprecision of cell-tower-based geolocation prevents us from distinguishing *presence at the protest*

---

<sup>15</sup>Appendix Table C.3 also reports the correlations between network centrality at one time step and network centrality at other time steps; we return to this below.

from *presence at a nearby home or business*. The data identify approximately 160,000 Partner Telecom subscribers who place a call in the vicinity of the protest during protest hours on September 1; when we drop those who live or work in parishes adjacent to the protest route, we are left with a sample of approximately 44,000 protesters.<sup>16</sup>

The use of two distinct *behaviorial* measures of political participation constitutes one of our contributions. First, behavioral measures allow us to avoid the desirability bias inherent in self-reported political participation—bias that could be stronger if social pressure indeed drives participation. Second, observing two different forms of participation (protest and petition signing) ensures that our results are not driven by specific features of one type of participation (like the fact that people often attend protests in groups). We discuss this advantage in detail below.

### 3.4 Matching participants to similar non-participants

The theoretical results summarized in Section 2 imply that, in equilibrium, people with greater communication centrality should be more likely to participate than otherwise identical people who are less exposed. The “otherwise identical people” part of these models poses an empirical challenge. In theoretical work, network position and other characteristics—like socioeconomic status—are assumed to be independent. In reality, they are correlated. We partially address this challenge by matching on observables.

For each protester in a random sample of 5,000 protesters, we identify a non-protester who is (a) registered to vote at the same polling place, (b) of the same political party, (c) likely to spend a similar amount of time in Caracas in a given month (excluding the protest day), (d) of the same gender, (e) on the same type of cell-phone plan (pre- or post-paid), (f) of a similar age, and (g) of a similar level of geographic mobility (as measured by the spatial variance, across days, of her cell phone).<sup>17</sup> Matching on polling place is perhaps the most important: polling places are small and neighborhoods are highly segregated (by class), so comparing people within polling place increases the likelihood of similarity along many dimensions. Relative to other studies of social network and political protests, we match participants and non-participants on a rich

---

<sup>16</sup>That we observe only 160,000 of the approximately 700,000 (estimated) protesters at the *Toma de Caracas* makes sense given (a) the fact that we only have cell-tower data for subscribers of our Partner Telecom, and (b) the fact that we only have cell-tower-routing data for phone calls, not for text messages.

<sup>17</sup>For details, see Appendix B. Results are similar when we also match on average user call frequency.

set of observable characteristics.<sup>18</sup>

## 4 Results: Communication Centrality and Political Participation

Descriptively, how does communication centrality relate to socio-economic characteristics and to political participation? We study these correlations using a representative sample of individuals from six large Venezuelan states. For these six states, we obtained census-tract maps that allow us to link neighborhood characteristics (like education) to electoral precincts (and thus to the voters in our data).<sup>19</sup> We re-weight the sample so that it resembles the overall Venezuelan population.

Communication centrality is weakly correlated with (neighborhood-level) poverty and education, but it is strongly correlated with both forms of political participation. Appendix Figure C.2 graphs the nonparametric bivariate relationships. To investigate how the (descriptive) relationship between communication centrality and participation changes when we account for other factors, we estimate:

$$\text{Participate}_i = \gamma_0 + \gamma_1 N_{i,5}^a + \gamma_2 \mathbf{X}_i + \eta_i \quad (1)$$

where  $N_{i,5}^a$  is communication centrality at  $t = 5$ , and  $\mathbf{X}_i$  includes (a) municipality fixed effects, (b) whether person  $i$  is registered with the United Socialist Party of Venezuela (PSUV, the party of the government), (c) the proportion of adults (over 25) in the neighborhood who have a college degree; and (d) the proportion of households in the neighborhood with a cement floor.

Table 1 reveals that communication centrality predicts larger changes in protest participation than does party registration or neighborhood socioeconomic characteristics (Columns 1–3; coefficients are scaled so as to capture the predicted change in participation associated with moving from the 5th to the 95th percentile of the independent variable). Communication centrality is also highly predictive of signing the petition. Moving from the 5th to the 95th percentile of communication centrality predicts a 4.5-percentage-point increase in the prob-

---

<sup>18</sup>Larson et al. (2019), for example, compared Twitter users who tweeted from the location of the Charlie Hebdo protests using a related hashtag to users who used related hashtags within Paris but at least 5km from the protest site. This strategy matched subjects on political interest and ability to attend, but not on socio-economic status or other characteristics.

<sup>19</sup>The six states are Aragua, Carabobo, the Federal District, Lara, Miranda, and Vargas.

Table 1: Communication centrality predicts political participation  
 Predicted change in participation associated with moving from the 5th to the 95th percentile of each independent variable (or from zero to one, for indicators). Based on estimates of Equation 1; robust standard errors in parentheses.

	Protest			Petition		
	(1)	(2)	(3)	(4)	(5)	(6)
Communication centrality ( $N_{i,5}^g$ )	0.027 (0.009)	0.040 (0.01)	0.040 (0.01)	0.046 (0.02)	0.042 (0.02)	0.041 (0.01)
Registered in PSUV (gov't party)		-0.020 (0.01)	-0.018 (0.01)		-0.113 (0.02)	-0.101 (0.01)
Signed petition		0.017 (0.006)	0.016 (0.006)			
% Neighborhood w/ cement floor			-0.009 (0.009)			-0.042 (0.05)
% Neighborhood w/ college degree			0.007 (0.008)			0.081 (0.05)
$N$	8033	8033	8033	8033	8033	8033
Municipality FEs		✓	✓		✓	✓

ability of signing the petition—about half of the difference associated with moving from the 5th to the 95th percentile of neighborhood education (c.f. Booth and Seligson, 2006).

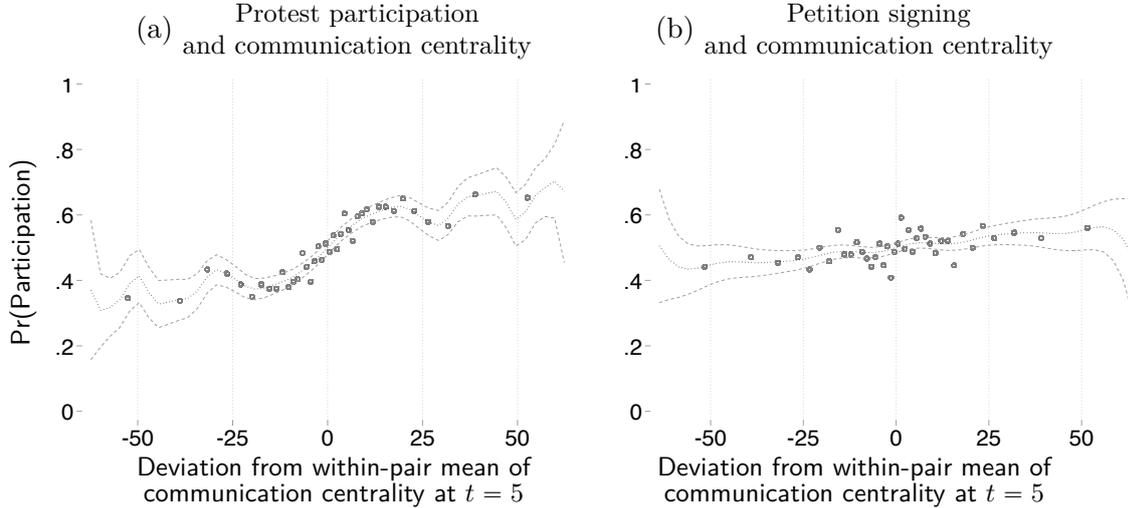
#### 4.1 Political participation also increases with communication centrality

These correlations are suggestive, but we can learn more by comparing the communication centrality of participants to that of observationally similar non-participants. The analysis in this section describes variation *within* matched pairs (Section 3.4), asking whether protesters and/or petition signers have higher or lower communication centrality than non-participants who share the same polling place, party registration, travel habits, and age group, among other characteristics (see Section 3.4 above for details). We use  $\overline{N_{pair}}$  to denote the within-pair mean of communication centrality, such that  $N_i - \overline{N_{pair}}$  denotes person  $i$ 's deviation from the within-pair mean of communication centrality. Except where otherwise noted, we use communication centrality measured at time step  $t = 5$ ; results for other time steps are similar.<sup>20</sup>

<sup>20</sup>In principle, it would be interesting to compare the effects of exposure at different time steps: does exposure at  $t = 2$  matter more than exposure at  $t = 5$ ? How much more? In practice, exposure is too highly correlated across time steps to allow us to make these distinctions, as Appendix Table C.3 reveals.

Figure 1: Communication centrality predicts political participation

Fig. (a) plots the probability that a person protests given her deviation from the within-matched-pair mean of communication centrality ( $N_{i,5} - \overline{N_{pair,5}}$ ). Figure (b) plots the probability that a person signs the recall referendum petition conditional on her deviation from the within-pair mean of communication centrality.



Lines shows fitted values (and 95% C.I.s) from local linear regression using a Gaussian kernel with the rule-of-thumb bandwidth from Fan and Gijbels (1996), 110–113. The regression is fit to the raw (unbinned) data. The points mark mean participation rates in bins of 2.4 percentiles (40 equally sized bins). For visual clarity, we exclude the top and bottom 2% (we do not trim the sample when estimating Equation 2).

Figure 1 plots the probability of protest participation and of petition signing against deviations from the within-pair mean of communication centrality,  $N_{i,5}^g - \overline{N_{pair,5}^g}$ . Even within matched pairs, the probability of political participation increases with communication centrality, consistent with the notion that communication centrality facilitates social sanctioning and thereby encourages protest participation. Petition signing rates also increase with communication centrality with pairs, though more slowly; we interpret the differences in Section 4.3 below.

Figure 1 reveals that the within-pair nonparametric relationship between communication centrality and political participation is close to linear. With this in mind, we estimate:

$$\text{Participate}_{ip} = \gamma_p + \beta N_{ip,5} + \epsilon_{ip} \quad (2)$$

where  $\text{Participate}_i$  is an indicator for whether person  $i$  in pair  $p$  participates (in the protest, in signing the petition, or in a placebo outcome, as indicated in each table);  $\gamma_p$  are fixed effects for each pair  $p$ ,  $N_{ip,5}$  is the communica-

Table 2: Communication Centrality Predicts Political Participation  
 Predicted change in probability of participation associated with a one-s.d. increase in communication centrality (one s.d. of the within-pair distribution of communication centrality), based on estimates of Eqn. 2.

	Protest (1)	Petition (2)
Communication centrality ( $N_i^g$ )	0.077 (0.01)	0.028 (0.010)
N	10K	10K

Standard errors, clustered by pair, in parentheses.

tion centrality of person  $i$  in pair  $p$  at  $t = 5$ ,<sup>21</sup> and  $\epsilon_{ip}$  is a person-specific shock term. This specification thus exploits *within-matched-pair* variation in communication centrality (Mummolo and Peterson, 2018).

Table 2 reports the estimates of Equation 2. A one-standard-deviation increase in within-pair communication centrality predicts a 7.7-percentage-point increase in protest participation, while a one-standard-deviation increase in within-pair communication centrality predicts a 2.8-percentage-point increase in the probability of signing the petition. Again, these correlations are consistent with the idea that communication centrality facilitates social sanctioning as a solution to the collective action problem.

One obvious concern is that the estimates in Table 2 might capture the effect of some unobserved characteristic that remains unbalanced across the two groups (and that is correlated with communication centrality). For example, though we match on observables correlated with socio-economic status (like neighborhood of residence), it might nevertheless be the case that protesters and petition-signers are richer or better educated than their matched non-participant counterparts.

To evaluate this possibility, we repeat our analysis with a placebo outcome that should be associated with socioeconomic status but not (necessarily) with communication centrality: presence in the area of Caracas where the protest took place, but on non-protest days (placebo dates). People who travel to Caracas for work or pleasure likely have more resources than those who do not. If our matching strategy failed to produce balance on socioeconomic status, we would expect “placebo protesters” (people making calls from the protest area on non-protest days) to have higher communication centrality

<sup>21</sup>As noted above, results are similar for other time steps.

than observationally similar people not traveling to Caracas.

To conduct this analysis, we first draw new samples of placebo protesters: people making calls from the protest area on eight non-protest Thursdays (the protest took place on a Thursday). For each placebo protester, we find an observationally similar person who was *not* in Caracas on that day, using the same matching strategy described in Section 3.4 above. We then measure communication centrality for placebo participants and non-participants, and estimate Equation 2 using this new matched sample.

Table 3: Communication centrality does *not* predict placebo activities

Predicted change in probability of participation associated with a one-standard-deviation increase in communication centrality (that is, one s.d. of the within-pair distribution of communication centrality), based on estimates of Eqn. 2.

	8/4	8/11	8/18	8/25	<b>9/1</b>	9/8	9/15	9/22	9/29
Communication centrality ( $N_i^g$ )	0.031 (0.010)	0.016 (0.01)	0.003 (0.010)	0.013 (0.010)	<b>0.077</b> <b>(0.01)</b>	0.006 (0.010)	0.010 (0.010)	-0.000 (0.010)	0.013 (0.010)
N	10K	10K	10K	10K	<b>10K</b>	10K	10K	10K	10K

Standard errors, clustered by matched pair, in parentheses.

If our matching strategy failed to capture important dimensions of socioeconomic status, we would expect people who travel to the protest area in Caracas on *any* date to have higher communication centrality than their matched counterparts. Instead, as the estimates in Table 3 reveal, communication centrality is much less predictive of the placebo outcome than it is of protest participation. One one-standard-deviation increase in communication centrality predicts a 1.1-percentage-point increase in the probability of presence in Caracas on placebo dates (compared to 7.7 percentage points for actual protest participation).<sup>22</sup>

## 4.2 Participation increases with exposure to *participants*

One concern with the analysis thus far is that we have focused on communication centrality, or the total number of peers who are likely to hear about a person’s behavior, whereas (in theory) only *participants* have the authority to sanction non-participation. Who are shirkers to criticize shirkers?

<sup>22</sup>In principle, we could conduct an analogous placebo-day analysis using exposure to participants (“number of participants informed in the information diffusion process”), rather than communication centrality. In practice, though, we would not expect these placebo-day estimates to be zero: assuming some homophilic tendency among people who travel to downtown Caracas on a given day, exposure to “participants” (scare quotes indicating placebo-day participants) *should* predict “participation.”

In this section, we show that the results are robust to focusing instead on exposure *to participants*—that is, the number of protest attendees or petition signers likely to hear about a person’s behavior.

Table 4 shows the relationship between participation and *exposure to participants*. Clearly, exposure to participants does predict participation—indeed, within matched pairs, it predicts an even larger change in participation rates than does communication centrality. Columns (2) and (4) report the estimates: a (within-pair) one-standard-deviation increase in exposure to participants predicts a 20.1-percentage-point increase in protest participation and a 4.9-percentage-point increase in the probability of signing the petition (compared to 7.7-percentage-points and 2.8-percentage-points, respectively, for one-standard-deviation changes in communication centrality).

Table 4: Exposure to Participants Predicts Participation  
 Predicted change in probability of participation associated with a one-std.-dev. increase in exposure to participants (that is, one s.d. of the within-pair distribution of network exposure), based on estimates of Eqn. 2.

	Protest		Petition	
	(1)	(2)	(3)	(4)
Communication centrality ( $N_i^g$ )	0.077 (0.01)		0.028 (0.010)	
Exposure to participants ( $N_i^p$ )		0.201 (0.01)		0.049 (0.010)
N	10K	10K	10K	10K

Standard errors, clustered by pair, in parentheses.

While these results are consistent with our claims, this is not our preferred specification. First, it is unclear whether the moral authority to sanction is restricted to participants (Banerjee et al., 2013); people unable to attend (for health or work reasons, for example) might nevertheless pressure others to participate. Likewise, participants and non-participants alike can provide social approbation as a reward for participating. And moreover, any tendency for participants to preferentially befriend other participants (a form of homophily) would inflate the correlation between exposure to participants and participation rates.

### 4.3 Protest vs. petition signing: what explains the difference?

The slopes in Figure 1 and the estimates in Tables 2 and 4 reveal that, within matched pairs, people with high communication centrality are more likely to participate than people with low communication centrality. But the relationship is stronger for protest participation than for petition signing. Why?

*Sociability* is almost certainly correlated both with communication centrality and with protest participation. A protest is an inherently social activity, usually attended in the company of friends, and thus protesters might have higher communication centrality than similar non-protesters for reasons unrelated to social sanctioning or, for that matter, to any solution to the collective action problem. In other words, we might observe the same pattern if we were to compare (say) people at a nightclub to observationally similar homebodies, even though clubbing is not subject to the collective action problem.

Petition signing, in contrast, eludes the sociability confound. Like protests, signature drives suffer from the collective action problem: participation is costly, and everyone (participants and non-participants alike) benefits from success. The temptation to free ride on others' petition signatures might be especially strong in Venezuela, where signatories of an earlier petition were fired from government jobs and otherwise discriminated against (Hsieh et al., 2011). But unlike protesting, petition signing is not a social event.

We interpret the difference in slopes between Figure 1a and 1b (and the corresponding difference in point estimates in Tables 2 and 4) as evidence that *sociability* drives part of the relationship between protest participation and communication centrality.

This finding would then affect the interpretation of related results in the literature. For example, Enikopolov et al. (2017a) find that an online social network increased protest participation in Russia, interpreting this result as evidence that the social network helped “solve the collective action problem” by “reducing the costs of coordination.” But if the social network also increased attendance at (say) concerts, we might interpret their results instead as evidence that the social network simply facilitated social activities, whether or not those activities were subject to the collective action problem. These two interpretations have different political implications: if the social network helped solve the collective action problem, we might expect that it would also enable other types of political activities; if it merely facilitates social events, we would not expect that it would affect other, less-social political activities, like voting. And indeed, Enikopolov et al. (2017a) find no evidence that the

social network increased anti-government votes.

Of course, there are other possible explanations for our finding that communication centrality predicts larger increases in protest participation than in petition signing. For example, protesting might be more *visible*: perhaps people are more likely to know whether their peers protested than to know whether their peers signed a petition. If so, the difference in slopes in Figure 1 could indeed be driven by a difference in the strength of social sanctioning, rather than by the sociability confound.

To evaluate this possibility, we included original questions on an in-person, nationally representative survey of Venezuelans conducted in September, 2018.<sup>23</sup> The results suggest that a person’s decision to protest is no more or less visible than her decision to sign the petition: the proportion of people who report knowing whether a friend participated is similar for both activities.<sup>24</sup>

Nor is protesting more likely to incur social sanction (or social approbation). In a separate question, we asked whether friends or family members “told [the respondent] what they thought of [the respondent’s] decision” to attend (or not attend) the protest, or whether “they kept their opinions to themselves.” We asked the same question about petition signing. If people were more likely to opine about their friends’ protest attendance (or non-attendance) than about their friends’ petition signatures (or lack thereof), we might interpret the difference in slopes (Figure 1a vs. Figure 1b) as evidence simply that social sanctioning is a more common tool for protest mobilization than for petition drives. Instead, our survey data suggest the opposite: more respondents reported hearing others’ opinions about their decision to sign (or not sign) the petition (44%) than about their decision to attend (or not attend) the protest (28%).

Our petition-signing results thus serve two purposes. First, they suggest that our results on protest activity—and similar results in the literature—may be driven not by the relationship between social networks and collective action *per se*, but rather (or also) by the relationship between social networks and social activities. At the same time, the non-zero relationship between communication centrality and petition signing suggests that sociability is not the whole story. Rather, the fact that communication centrality predicts petition signing at all is consistent with the idea that communication centrality facilitates social

---

<sup>23</sup>The firm Datanálisis included our questions on their regular quarterly survey.

<sup>24</sup>We asked each respondent both about protesting and about petition-signing, randomizing the order of the questions. Results are nearly identical when we restrict the analysis to the question that respondents answered first. Appendix D presents the full survey instrument, in English and Spanish.

sanctioning and thereby discourages free riding.<sup>25</sup>

## 5 Discussion

Until recently, it was impossible to observe social networks at scale. That rendered unknowable many facts about the role of social ties in solving collective action problems. In particular, the theoretical prediction that communication centrality—a measure of how quickly others hear about a person—facilitates social sanctioning, and thereby political participation, has gone untested.

In this paper, we use newly available data to study the relationship between communication centrality and political participation. We map social networks for nearly the entire country of Venezuela using cell-phone meta-data. We then estimate communication centrality and pair it with behavioral measures of two different forms of political participation: protest and petition signing.

Consistent with theory, we find that communication centrality predicts political participation. This is true even within a matched sample of participants and non-participants who live in the same small neighborhood and share the same party registration and gender, among other characteristics. In other words, we are able to compare participants with non-participants who appear similar across more observables than have been used in past work. Moreover, these participants and matched non-participants engage in placebo activities that should be associated with socio-economic status—but that are not subject to the collective action problem—at nearly identical rates. And again, even within these matched pairs, communication centrality strongly predicts political participation.

Together with qualitative information and original survey data, we argue that communication centrality and political participation are connected through social sanctioning. Specifically, we rule out alternative explanations like the possibility that communication centrality simply facilitates *awareness* of political activities. We therefore argue not only that “networks matter” for political

---

<sup>25</sup>One potential caveat is that perhaps *extroversion* drives political participation in general (Gerber et al., 2011), even forms of participation (like petition-signing) that do not occur in large groups. But studies of personality type and participation find that extroversion is correlated only with “the tendency to engage in those forms of political participation that involve social interaction, especially interaction in large groups, but the influence of extroversion on political participation dissipates when focus turns to more individualistic behaviors” (Mondak 2010, 159–160). In our context, petition signing took place over multiple days at hundreds of locations; while it undoubtedly involved more social interaction than (say) putting up a yard sign, it did not typically involve “interaction in large groups.” For that reason, extroversion strikes us as an unlikely explanation for the observed relationship between petition signing and communication centrality.

participation, but specifically that they facilitate social sanctioning (and social approbation) as a solution to the collective action problem. While some aspects of the Venezuelan case are unusual—for example, few protests are as large as the one studied in this paper—we nevertheless expect that social sanctioning facilitates collective action elsewhere, too.

Our results also imply a word of caution about the study of social networks and collective action more generally. Much of the empirical work on this subject has focused on protests, for the good reason that they are politically important, and for the convenient reason that cell-phone location allows researchers to observe participation (Steinert-Threlkeld, 2017; Larson et al., 2019; Enikopolov et al., 2017a). But in comparing protest to another political activity—petition signing—we find that network position may be correlated with protest activity for reasons that have nothing to do with collective action, coordination, awareness, or other reasons commonly cited in the literature. Protest is an inherently social activity, more like attending a street fair than like casting a ballot. Extroversion or sociability may therefore drive correlations between protest and a person’s place in the social network. Studying more solitary political activities strikes as a worthy goal for future work.

## References

- Albertus, M. (2015). The role of subnational politicians in distributive politics: political bias in venezuela’s land reform under Chávez. *Comparative Political Studies*.
- Arias, E., Balán, P., Larreguy, H., and Marshall, J. (2019). Information provision, voter coordination, and electoral accountability: Evidence from mexican social networks. *American Political Science Review*.
- Auyero, J. (2000). *Poor People’s Politics: Peronist Survival Networks and the Legacy of Evita*. Duke University Press.
- Aytaç, S. E. and Stokes, S. C. (2019). *Why Bother?: Rethinking Participation in Elections and Protests*. Cambridge University Press.
- Banerjee, A., Chandrasekhar, A., Duflo, E., and Jackson, M. O. (2014). Gossip: Identifying Central Individuals in a Social Network. *SSRN Working Paper Series*, 20422.
- Banerjee, A., Chandrasekhar, A. G., Duflo, E., and Jackson, M. O. (2013). The diffusion of microfinance. *Science*.
- Barthélemy, M. (2011). Spatial Networks. *Physics Reports*, 499(1-3):1–101.
- Björkegren, D. (2018). The Adoption of Network Goods: Evidence from the Spread of Mobile Phones in Rwanda. *The Review of Economic Studies*.
- Blondel, V., Krings, G., and Thomas, I. (2010). Regions and Borders of Mobile Telephony in Belgium and in the Brussels Metropolitan Zone. *Brussels Studies*.
- Blumenstock, J., Cadamuro, G., and On, R. (2015). Predicting Poverty and Wealth from Mobile Phone Metadata. *Science*, 350(6264):1073–1076.
- Blumenstock, J. E. (2012). Inferring Patterns of Internal Migration from Mobile Phone Call Records: Evidence from Rwanda. *Information Technology for Development*, 18(2):107–125.
- Booth, J. A. and Seligson, M. A. (2006). Inequality and democracy in latin america: Individual and contextual effects of wealth on political participation. *Duke University Workshop on Poverty and Democracy*.
- Calvo, E. and Murillo, M. V. (2013). When parties meet voters: Assessing political linkages through partisan networks and distributive expectations in argentina and chile. *Comparative Political Studies*.

- Cantoni, D., Yang, D. Y., Yuchtman, N., and Zhang, Y. J. (2019). Protests as strategic games: Experimental evidence from hong kong’s antiauthoritarian movement. *The Quarterly Journal of Economics*.
- Centola, D. (2018). *How Behavior Spreads: The Science of Complex Contagions*. Princeton University Press, Princeton.
- Centola, D. and Macy, M. (2007). Complex Contagions and the Weakness of Long Ties. *American Journal of Sociology*, 113(3):702–734.
- Christensen, D. and Garfias, F. (2018). Can You Hear Me Now?: How Communication Technology Affects Protest and Repression. *Quarterly Journal of Political Science*.
- Datanalisis (2016). Encuesta omnibus. *Primer Trimestre*.
- Dionne, K. Y. (2015). Social Networks, Ethnic Diversity, and Cooperative Behavior in Rural Malawi. *Journal of Theoretical Politics*, 27(4):522–543.
- Eagle, N., Pentland, S., , and Lazer, D. (2009). *Inferring friendship network structure by using mobile phone data*. Proceedings of the National Academy of Sciences.
- Enikopolov, R., Makarin, A., and Petrova, M. (2017a). Social media and protest participation: Evidence from russia. *Working Paper*, page <https://ssrn.com/abstract=2696236>.
- Enikopolov, R., Makarin, A., Petrova, M., and Polishchuk, L. (2017b). Social image, networks, and protest participation. *Working Paper*.
- Eubank, N., Grossman, G., Platas, M., and Rodden, J. (2018). Social Networks and Turnout in Low Salience Elections: Evidence from Uganda. *Working Paper*, page <https://ssrn.com/abstract=3043663>.
- Fafchamps, M. and Vicente, P. C. (2013). Political Violence and Social Networks: Experimental Evidence from a Nigerian Election. *Journal of Development Economics*, 101:27–48.
- Fan, J. and Gijbels, I. (1996). *Local Polynomial Modelling and its Applications*. CRC Press.
- Fearon, J. D. and Laitin, D. (1996). Explaining Interethnic Cooperation. *American Political Science Review*, 90(4):715–735.

- Ferrali, R., Grossman, G., Platas, M., and Rodden, J. (2019). It takes a village: Peer effects and externalities in technology adoption. *American Journal of Political Science*.
- Gerber, A. S., Green, D. P., and Larimer, C. W. (2008). Social Pressure and Voter Turnout: Evidence from a Large-Scale Field Experiment. *American Political Science Review*, 102(1):33–48.
- Gerber, A. S., HUBER, G. A., Doherty, D., and Dowling, C. M. (2011). The Big Five Personality Traits in the Political Arena. *Annual Review of Political Science*, 14(1):265–287.
- Gerber, A. S., Huber, G. A., Doherty, D., and Dowling, C. M. (2016). Why People Vote: Estimating the Social Returns to Voting. *British Journal of Political Science*, 46(2):241–264.
- Gonzalez, F. (2017). Collective action in networks: Evidence from the chilean student movement. *Working Paper*, page <https://ssrn.com/abstract=2943760>.
- Granovetter, M. S. (1973). The Strength of Weak Ties. *American Journal of Sociology*, 78(6):1360–1380.
- Hsieh, C.-T., Miguel, E., Ortega, D., and Rodriguez, F. (2011). The price of political opposition: Evidence from venezuela’s maisanta. *American Economic Journal: Applied Economics*.
- Jackson, M. O. (2008). *Social and Economic Networks*. Princeton University Press.
- Jackson, M. O. (2019). A typology of social capital and associated network measures. *Social Choice and Welfare*.
- Jackson, M. O., Rodriguez-Barraquer, T., and Tan, X. (2012). Social capital and social quilts: Network patterns of favor exchange. *American Economic Review*.
- Jackson, M. O., Rogers, B. W., and Zenou, Y. (2017). The economic consequences of social-network structure. *Journal of Economic Literature*.
- Kandori, M. (1992). Social Norms and Community Enforcement. *The Review of Economic Studies*, 59(1):63–80.
- Kronick, D. (2015). Why would a president wreck the economy before an election? *The Washington Post*.

- Kuran, T. (1991). Now out of never: The element of surprise in the East European revolution of 1989. *World Politics*, 44(01):7–48.
- Larson, J. (2017a). Networks and interethnic cooperation. *Journal of Politics*.
- Larson, J. and Lewis, J. (2020). Measuring networks in the field. *Political Science Research and Methods*.
- Larson, J. M. (2017b). Why the West Became Wild. *World Politics*, 69(4):713–749.
- Larson, J. M. and Lewis, J. I. (2017). Ethnic Networks. *American Journal of Political Science*, 61(2):350–364.
- Larson, J. M., Nagler, J., Ronen, J., and Tucker, J. A. (2019). Social Networks and Protest Participation: Evidence from 130 Million Twitter Users . *Forthcoming, American Journal of Political Science*.
- Little, A. T. (2015). Communication Technology and Protest. *Journal of Politics*, 78(1):152–166.
- Miritello, G., Moro, E., Lara, R., Martínez-López, R., Belchamber, J., Roberts, S. G. B., and Dunbar, R. I. M. (2013). Time as a limited resource: Communication strategy in mobile phone networks. *Social Networks*, 35(1):89–95.
- Mondak, J. J. (2010). *Personality and the Foundations of Political Behavior*. Cambridge University Press, Cambridge, MA.
- Mummolo, J. and Peterson, E. (2018). Improving the interpretation of fixed effects regression results. *Political Science Research and Methods*.
- Newman, M., Barabási, A.-L., and Watts, D. J. (2006). ***The Structure and Dynamics of Networks***. Princeton University Press.
- Nickerson, D. (2008). Is voting contagious? evidence from two field experiments. *American Political Science Review*.
- Noticia al Día (2016). Así se desarrolla en maracaibo el segundo día de recolección de firmas para el revocatorio. <https://bit.ly/2M8bFwP>.
- Olson, M. (1965). *The Logic of Collective Action*. Harvard University Press.
- Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press.

- Paler, L., Marshall, L., and Atallah, S. (2018). The social costs of public political participation: Evidence from a petition experiment in Lebanon. *The Journal of Politics*.
- Panagopoulos, C. (2010). Affect, Social Pressure and Prosocial Motivation: Field Experimental Evidence of the Mobilizing Effects of Pride, Shame and Publicizing Voting Behavior. *Political Behavior*, 32(3):369–386.
- Rodriguez, F. (2016). Chronicle of a march foretold. *Torino Capital Venezuela this Week*.
- Rodriguez, F. (2017). As U.S. sanctions Maduro and hints at regime change, a debate on resolving the crisis in Venezuela. *Democracy Now*.
- Rojo, G. and Wibbels, E. (2014). Clientelism, Public Services and Elections in the Slums of Jaipur and Udaipur. *Presented at Stanford Comparative Politics Seminar*.
- Rolfe, M. (2012). *Voter Turnout: A Social Theory of Political Participation*. Cambridge University Press, Cambridge.
- Siegel, D. A. (2009). Social Networks and Collective Action. *American Journal of Political Science*, 53(1):122–138.
- Sinclair, B. (2012). *The Social Citizen: Peer Networks and Political Behavior*. The University of Chicago Press, Chicago.
- Steinert-Threlkeld, Z. C. (2017). Spontaneous Collective Action: Peripheral Mobilization During the Arab Spring. *American Political Science Review*, 111(2):379–403.
- Steinert-Threlkeld, Z. C. (2018). Activism and Protest in Authoritarian Regimes. *Working Paper*, page [shorturl.at/jksvz](https://shorturl.at/jksvz).
- Steinert-Threlkeld, Z. C., Mocanu, D., Vespignani, A., and Fowler, J. H. (2015). Online social networks and offline protest. *EPJ Data Science*, 4(1).
- Stokes, S. C., Dunning, T., Nazareno, M., and Brusco, V. (2013). *Voters, Brokers, and Clientelism: The Puzzle of Distributive Politics*. Cambridge University Press.
- Toro, F. (2014). Rash repression in Venezuela. *The New York Times*.
- Toro, F. (2016). Signed, sealed, and delivered. *Caracas Chronicles*.

- Ujhelyi, G. and Calvo, E. (2011). Political screening: Theory and evidence from the argentine public sector. *Working Paper*.
- Wesolowski, A., Eagle, N., Tatem, A. J., Smith, D. L., Noor, A. M., Snow, R. W., and Buckee, C. O. (2012). Quantifying the Impact of Human Mobility on Malaria. *Science*, 338(6104):267–270.
- Wolitzky, A. (2012). Cooperation with network monitoring. *Review of Economic Studies*.