

Social Networks and the Mass Media

DAVID A. SIEGEL *Duke University*

How do global sources of information such as mass media outlets, state propaganda, NGOs, and national party leadership affect aggregate behavior? Prior work on this question has insufficiently considered the complex interaction between social network and mass media influences on individual behavior. By explicitly modeling this interaction, I show that social network structure conditions media's impact. Empirical studies of media effects that fail to consider this risk bias. Further, social network interactions can amplify media bias, leading to large swings in aggregate behavior made more severe when individuals can select into media matching their preferences. Countervailing media outlets and social elites with unified preferences can mitigate the effect of bias; however, media outlets promulgating antistatus quo bias have an advantage. Theoretical results such as these generate numerous testable hypotheses; I provide guidelines for deriving and testing hypotheses from the model and discuss several such hypotheses.

Do the media affect opinion leading up to the Iowa caucuses the same way they do in the California primary? Under what conditions can national party leaders sway state politicians' votes? When will the Voice of America or a pro-democracy NGO or someone broadcasting over Facebook be effective in spurring democratic movements like the Arab Spring? Can media pundits drive opinion, or do they merely preach to the choir? Each of these is an example of a broader question: How do the mass media¹ affect aggregate opinion and behavior?

To begin to answer this question, I develop a novel theory of aggregate opinion and behavior. The theory considers a heterogeneous population of individuals who must choose between dichotomous options. It incorporates the interaction of social network and mass media influences at the individual level; its key assumption is that the more others choose an option, the more one is apt to do so as well. In the theory, social networks provide information about the choices of those to whom one is directly connected, while the mass media provide (potentially biased) information about aggregate choice. The theory thus applies to, for example, voter turnout and political participation (e.g., Gerber, Green, and Larimer 2008; Lake and Huckfeldt 1998; Leighley 1990; McClurg 2003; Rolfe 2012), opinion formation (e.g., Beck et al. 2002; Druckman and Nelson

2003; Huckfeldt and Sprague 1995), protests and social movements (e.g., Kuran 1991; McAdam 1986), and vote choice (e.g., Beck 2002; Huckfeldt and Sprague 1995; Ryan 2011; Sinclair 2012; Sokhey and McClurg 2012).

Three major results follow from this theory. All hold both when individuals treat media identically and when they select into media in line with their preferences. First, understanding the aggregate effect of the media generally requires considering social networks, because *social network structure conditions media's impact*. For example, additional weak ties between disparate social groups can reduce the media's impact, and the presence of unified social elites can eliminate the media's impact entirely in the aggregate. *Empirical studies of media impact that fail to consider media's interaction with social networks risk bias*.

Second, *social networks can amplify the effect of media bias*. A biased media outlet that systematically under- or over-reports a poll of the population by a only a few percentage points can in some cases swing aggregate behavior (e.g., turnout or vote share) by over 20% in either direction due to positive feedback within the network. Open advocates in the media can have a yet larger impact even when not comparatively influential. Unified social elites limit the effect of media bias, but cannot fully counter an advocate; selection into media, made ever easier with technological improvements, tends to *enhance* the effect of bias. We should therefore expect media bias to become increasingly important to aggregate behavior.

Third, multiple biased media outlets promulgating countervailing messages do not necessarily produce neutral aggregate outcomes. *Media bias is generally more effective at driving the population away from a status quo option than towards it*. Social networks are the cause of this; selection into media again exacerbates this trend. Thus, newer, antiestablishment media outlets can have a greater aggregate impact than older outlets that are biased toward the status quo. This suggests the preferential focus and growth of outlets opposing the status quo if viewers and funding follow media impact.

General theoretical results like these generate numerous testable hypotheses. For example, the media

The author thanks Jason Barabas, Jon Bendor, Ted Carmines, Jamie Druckman, John Freeman, Matt Golder, Sona Golder, Bob Jackson, Jenn Jerit, Kris Kanthak, Özge Kemahlioglu, Charlotte Lee, Valerie Martinez-Ebers, Adam Meirowitz, Scott McClurg, Will Moore, Chris Reenock, John Ryan, John Scholz, Jake Shapiro, Anand Sokhey, Jeff Staton, Jim Stimson, Craig Volden, Jon Woon, four very helpful anonymous reviewers, and audiences in the Political Economics group at the Stanford GSB, Political Science departments at FSU, GWU, Minnesota, Pittsburgh, and Stony Brook, and the Frank Batten School of Leadership and Public Policy at UVa. Any errors are my own.

David A. Siegel is Associate Professor, Department of Political Science, Duke University, Durham, NC 27708 (david.siegel@duke.edu); web: <http://people.duke.edu/~das76>.

¹ Mass media outlets in this article may be conventional sources of information such as the *New York Times* and *FOX News*, or other sources external to social networks such as state propaganda, NGOs, national party leadership, or Internet personalities. I elaborate on this in the next section.

are more effective at moving public opinion away from the status quo the poorer the social network is at spreading behavior. The theory's network typology provides a way to assess with minimal data how well different networks spread behavior. Applying it suggests that the media will have a relatively larger impact in caucus states if the caucus system encourages less communication across precincts than does the primary system.

Another example concerns elite networks, such as those often characterizing long-standing organizations. When can external media, biased relative to the network, effect change? The theory suggests that unified social elite preference much reduces media's impact. This has broad implications. For instance, the level of religious organizations' accommodation of gay marriage or the teaching of evolution under media pressure is likely to be inversely correlated with the degree of common preference of their leadership. The same is true for pressured changes in longstanding policies in organizations like the military or political parties. In each case significant media impact only occurs once there is variation in elite opinion, though all elites need not support change for the media to abet a global shift in opinion.

One final example is on an international scale. How will decreased state ability to limit access to external messages from media and NGOs alter authoritarian states' strategies in maintaining power? States opening up economically, like China, will be forced by technology to turn to strategies other than simply shutting off contact with the outside world, as in North Korea. The theory suggests that fruitful strategies to inhibit the spread of antiregime collective action involve promulgating states' own status-quo bias among elites while using superior local knowledge to take advantage of bottlenecks in social networks. This appears to be in line with what the Chinese government presently does (King, Pan, and Roberts 2012).

I introduce the theory and the formal computational model that instantiates it in the second and third sections, respectively. The fourth, fifth, and sixth sections present the model's results corresponding to unbiased media, biased media, and selection into media, respectively. The staged presentation is designed to help develop clear characterizations of the joint effects of networks and media. The seventh section illustrates how to derive and test hypotheses like those above; it also summarizes the model's results. The eighth section concludes with a discussion of model extensions. Additional results, model description, and detail on measurement may be found in an Online Appendix (<http://dx.doi.org/10.1017/S0003055413000452>).

AN INDIVIDUAL-LEVEL THEORY OF AGGREGATE BEHAVIOR

Though I present a theory of aggregate behavior, it is based on individual-level assumptions informed by what we know about the way personal characteristics,

social networks, and mass media outlets affect individual behavior. Due to this, the theory can explore the effect that *interactions* between these three factors have on aggregate behavior. As importantly, the theory incorporates empirically realistic heterogeneity across people in all three factors.

Prior research at the individual level identifies that personal characteristics such as education, age, and income correlate with political participation (e.g., Leighley 1990; McClurg 2003) and are related to political knowledge (e.g., Delli Carpini and Keeter 1996). Knowledge affects policy preferences (Althaus 2003) and vote choice (Bartels 1996; cf. Lau and Redlawsk 1997; Lupia and McCubbins 1998). The beliefs and behavior of others in one's social network alter one's own beliefs and behavior via both direct influence and the provision of information (e.g., Beck 2002; Chong 1991; Huckfeldt 2001; Huckfeldt and Sprague 1995; Leighley 1990; Klobstad 2011; Klobstad, Sokhey, and McClurg 2013; McAdam 1986; McClurg 2003, 2006; Mutz 2002; Petersen 2001; Ryan 2011).

Additionally, people are exposed to individuals, groups, and organizations external to one's network, such as mass media outlets, state propaganda, national party leaders, NGOs, and Internet personalities. These outlets can provide information, increasing political knowledge (e.g., Barabas and Jerit 2009; Delli Carpini and Keeter 1996; Graber 2004; Jerit, Barabas, and Bolsen 2006; Mondak 1995b; Zaller 1992) or informing individuals as to the preferences of influential elites with regard to political participation, opinion, and vote choice (e.g., Ansolabehere, Behr, and Iyengar 1991; Chiang and Knight 2008; Della Vigna and Kaplan 2007; George and Waldfogel 2006; Gentzkow 2006; Gerber, Karlan, and Bergan 2009; Ladd and Lenz 2008; Zaller 1992).

As this small sampling of large literatures indicates, individuals' decisions are influenced by the information they obtain via both local social networks and global media outlets. However, comparatively little scholarship has explored the three-way interaction of personal characteristics, social networks, and media (e.g., Ahn, Huckfeldt, and Ryan 2010; Beck et al. 2002; Berelson, Lazarsfeld, and McPhee 1954; Eveland et al. 2005; Druckman and Nelson 2003; Druckman 2004; Mondak 1995a; Mutz and Martin 2001; Rojas, Shah, and Friedland 2011).² Importantly, this scholarship highlights the powerful role of social interactions, which have been shown, for example, to be more influential than the media in affecting one's vote (Beck et al. 2002).

Even the literature examining this three-way interaction, though, is largely confined to the individual level. Simple aggregation to obtain insights on aggregate behavior is likely to be insufficient in light of research indicating that the structure of social networks interacts with the distribution of personal characteristics in

² A related literature in communications looks at the relationship between the media and social capital specified more broadly in terms of associational memberships and feeling towards neighbors (Beaudoin and Thorson 2004) or community integration (McLeod et al. 1996).

a complex, nonlinear fashion (e.g., Centola and Macy 2007; Siegel 2009). There is little reason to believe that the addition of the media simplifies this interaction, particularly since the media can indirectly affect one's behavior via their effect on the opinions and behavior of others within one's social network.

The theory I present allows a better understanding of the media's impact in the aggregate while still incorporating factors identified by prior research as affecting individual-level behavior. This requires accommodating the complex interactions between the effects of social network structure, media outlets, and the population's distribution of personal characteristics. To facilitate this, the computational formal model that instantiates the theory generalizes an existing model of social network structure (Siegel 2009) via the addition of one or more mass media outlets. Building off this earlier model enables the use of its results regarding the interaction between personal characteristics and social network structure, and thus allows us to focus on the media's impact.

Individuals in the model are defined by personal characteristics that make them intrinsically more or less motivated to support one of two options. They are influenced by two different sources of information that affect their decisions: social networks and the media. Social networks provide information about *local* support for each option. The structure of the network determines whose decisions each person finds influential. In other words, if the network specifies a direct connection between two people (i.e., each person is in the (immediate) social network of the other), then each takes the other's behavior into account in her decision-making. I elaborate on this in the following section.

A mass media outlet in the model, in contrast, is a *global* source of information. It can be any information source that (i) influences individual behavior, (ii) is external to the social network, and (iii) is globally accessible to all. Thus, though I use the phrases *media* and *media outlet* exclusively below for simplicity, the "media" can be any form of commercial or state-run media, state propaganda, NGO, national party leadership, Internet personality, or other information source that satisfies these three requirements.³

The media are defined by two characteristics: their *strength* and *bias*. *Strength* refers to the relative level of influence the media have compared to one's social network. It is related to concepts of attention and trust. For example, the media may have little influence if their information is ignored or if people do not trust their veracity or underlying motives.⁴ Strength is not identical to these concepts, however, as it also encodes the degree to which individuals pay attention to or trust the information arising from their social networks.

³ In the sixth section of the Online Appendix I weaken the assumption of global accessibility, so that not all have access to the media. Major results continue to hold, and I find additionally that in some cases less access to the media can increase aggregate support for one of the options.

⁴ The link between strength and trust allows analysis of the aggregate effect of changes in media trust over time (Cohen 2004), though I do not consider this question here.

For example, when considering support for a complex policy, even a relatively poorly trusted media might be more influential in one's choice than social connections lacking any policy expertise. Because media strength specifies relative influence, greater media strength necessarily means one's local social network will be relatively less influential. However, the model still allows cases in which media and social networks both have relatively little influence (when personal characteristics imply intrinsic motivation is very low or extremely high) or great influence (when intrinsic motivation is more middling).

The baseline media of the model accurately report information on both options by passing along the aggregate level of support for each option in the population. The most straightforward interpretation of the media's information is that it is the result of a poll of the population's support for each option. This type of media information is commonly observed in public opinion polls, polls of candidate support, and media reports of voter turnout and participation in protests, and it is growing increasingly available in industrialized nations. However, the media's information need not be this reductionist. Instead, the baseline media may fairly report more detailed facts and opinions about the options. Individuals draw inferences about each option from this fair reporting, and they use aggregate support for the options as a proxy for this inference.⁵ Either way the baseline media are purely informative, supplementing local information received from one's social network with global information on all individuals in the population.

Bias specifies the deviation between the information a media outlet reports and the true aggregate level of support for each option. I consider two types of bias. In the first type of bias, which I call *polling bias*, the media outlet systematically biases its report of the outcome of its poll by some amount (or it simply takes a biased poll). Though I do not model the choice of bias as strategic, a media outlet might have incentive to signal stronger (or weaker) support for the option it favors. Political parties' in-house pollsters are examples of this, as the desire to curry favor with their employers may lead them to bias their reports in their employers' favor.

The model also captures less straightforward bias when the media's information comprises more complex reporting than a simple poll. Bias could arise from the media altering their balance of stories due to editorial slant (Druckman and Parkin 2005) or a tendency to focus on the more interesting or sensational story (Groeling and Kernell 1998).⁶ Or bias could reflect divergence between media cues related to issue evolution (Carmines and Stimson 1990) and a present position of the population on the options that is colored by poor

⁵ Others' support is based in part on information that is private to them or available within their social networks. Aggregate support is thus informative about the options in a manner similar to the way support within one's social network is. Global and local sources of information will in general be differentially influential, though, as specified by the media strength parameter.

⁶ I thank Kris Kanthak for the latter insight.

information or poor comprehension of the alignment between the options and true preferences.⁷

In the second type of bias, which I call *advocacy*, the media outlet simply states a preference for one of the options, providing no information about aggregate support. The goal in advocacy is to sway the population toward one or the other option. As before, many goals could underlie advocacy beyond just the support of a biased media outlet's preferences. Advocacy represents the editorial power of the media or the influence of an external actor; it is a "one-message" model (Zaller 1992).

In the fourth section below I analyze a single, baseline media outlet and vary its strength. Because the baseline media provide unbiased information, results from this analysis generalize directly to multiple unbiased media outlets.⁸ In the fifth section I analyze media bias. There I first consider one media outlet, which might correspond to state-run media in an authoritarian regime, and then two media outlets, which might correspond to competitive state or private media outlets in democratic regimes. Results from the analysis of two biased media outlets generalize to multiple biased media outlets.

The analysis in the fourth and fifth sections assumes that media strength and bias vary independently. However, those intrinsically motivated to support one of the options might be more strongly influenced by media whose bias matches their proclivities, and less strongly influenced by media whose bias runs counter to their proclivities. In other words, bias and strength might be correlated. In the sixth section I account for this possibility by extending the model to allow individuals to select into different levels of media strength based on the bias of the media.

I focus my analysis in all three sections on the case in which one of the two options is the status quo, and all individuals begin supporting it. For political participation and social movements, the status quo is not participating. For opinion formation and vote choice, the status quo is an existing option such as a policy in place or an incumbent politician, as contrasted with an alternative such as a newly proposed policy or a challenging politician. For simplicity I subsequently call *participation* the option that is not the status quo; this should be read as "participation in support of" the option that is not the status quo in contexts other than political participation or social movements.

⁷ There is a broad literature on media bias (e.g., Niven 2002; Page 1996) but no definitive approach as to how to define it (Druckman and Parkin 2005, Niven 2002). One issue with any definition is specifying the baseline of no bias (Druckman and Parkin 2005). This model has a well-defined baseline (aggregate support) and so elides this issue. It is also general enough to cover many potential definitions; I have noted a few in the text. It is sufficient for application of the model that individuals respond appropriately to bias, as they seem to do (Druckman and Parkin 2005, Kahn and Kenney 2002).

⁸ Empirically the strength of this joint media message will likely increase with confirming evidence; however, since I vary media strength across its entire range this does not affect the analysis.

MODEL

There are three components of the computational formal model. The first is individual behavior, comprising the manner personal characteristics and social network and media influences interact to affect behavior. The second is social network structure. The third is the number, strength, and bias of media outlets. In this section I briefly discuss each. Far more detail, along with additional justification for modeling choices, can be found in the Online Appendix.⁹

Individual Behavior

The focus of the model is a finite population of individuals who each must decide among dichotomous options, one of which is the status quo. Support for the alternative to the status quo is *participation*. Each individual is defined by internal and external motivations to support the alternative to the status quo—in other words, motivations to participate. *Internal motivations* encompass all motivations deriving from personal characteristics that *do not depend on the participation of others*. Under the assumption that there are many such characteristics that affect individuals' decision-making, each distributed across the population in some unknown fashion, a central limit theorem suggests that the net of these will be drawn from a normal distribution. All individuals' internal motivations are drawn at random from this distribution.

External motivations cover everything that internal motivations do not—specifically, all factors that *do* depend on the participation of others. In the model, these factors arise from two sources of information: one's social network and the media. The information provided by social networks is the local participation rate (i.e., rate of support for participation) among those to whom one is connected.¹⁰ The information provided by the baseline media is the global participation rate. Biased media alter this information, as described below.

Global and local sources of information together dictate external motivations: external motivations are a linear combination of these two information sources. The media strength parameter, M_S , determines the weight on each of the information sources in the linear combination. The greater the parameter, the more weight on global information. A media strength of 1 implies that local information has no influence on one's decisions, while a media strength of 0 implies that global information has no influence on one's decisions. This modeling choice allows an easy connection to the less complex model of Siegel (2009) and provides a

⁹ The Online Appendix can be found at <http://people.duke.edu/~das76/Research.html>.

¹⁰ Networks are defined as the set of people whose participation directly influences one's decisions via observation of their choice of option to support. The logic of the model holds as well when individuals respond only to others' influence, and do not directly observe participation, as long as those doing the influencing participate themselves when the time comes. See Gerber, Green, and Larimer (2008) for a similar take on the role of internal and external incentives in compliance with a norm.

straightforward way to vary the importance of the media's information.

Since the combined information from local and global sources increases in others' participation, so too do external motivations. This may be due, for example, to supportive information or influence (e.g., Huckfeldt and Sprague 1995), increased safety-in-numbers (e.g., Kuran 1991), increasing shame from violating norms of fairness by shirking (e.g., Gould 1993), or an increasing chance of being punished by the group for deviations from a norm.

Individuals simultaneously decide whether or not to participate over a series of periods; these end only when no one wants to change one's decision. Each period's decision is a simple comparison of internal and external motivations: if the sum of the two motivations exceeds a cutoff, an individual participates in that period. Consequently, the greater both local and global participation rates are, the more likely one is to participate. Similarly, increased internal motivations also make one more likely to participate.

The first period begins with a random realization of individuals' internal motivations and all individuals not participating. Early participants are driven to participate by high internal motivations. Subsequent participants are drawn in by a combination of their internal motivations and others' participation. Some will never participate under any circumstances, perhaps because they derive a great deal of power, prestige, or money from the status quo. Though individuals are free to stop participating at any time, in practice they never do so since external motivations are monotonic in others' participation.¹¹ This implies that aggregate participation weakly increases with period. At some point no new people begin to participate and the model reaches its steady-state outcome. Analysis focuses on this steady state. In particular, I focus on the steady-state, or equilibrium, aggregate level of participation in the population. This is aggregate support for the non-status-quo option.

For each realization of individual motivations one of two steady-state outcomes can occur: either there is a cascade that results in nearly complete participation, or there is not a cascade and only a lower level of participation is reached. For a given population size and given internal motivation distribution parameters, one of three classes of outcome can occur in the model (Siegel 2009). Since these are related to the distribution of internal motivations, I call each a *motivation class*. First, there can be frequent cascades; this occurs in the *strong motivation class*. Second, there can be few cascades; this occurs in the *weak motivation class*. Third, cascades can happen in some realizations but not others; this occurs in the *intermediate motivation class*. More frequent cascades translate directly to greater average aggregate participation.

Within each motivation class, average aggregate participation is similar. More important, the impact of

network structure and the media on aggregate participation is the same in each class as well. This means that analysis can focus on differences across classes. In performing this analysis, I find that the effect of the interaction between network structure and the media is identical in form in both the weak and intermediate motivation classes, but differs in the strong class. I thus only vary internal motivations across the strong and the intermediate motivation classes in the results that follow, but all results for the latter class also apply directly to the weak class.¹²

Social Networks

Social networks detail the set of other individuals who *directly* influence one's behavior. While there are a tremendous number of potential networks one could explore, I focus on a typology of qualitative network structures that mirror commonly observed empirical networks (Siegel 2009). I choose this focus to help make the analysis broadly useful: comparatively little data are necessary to discern network type, allowing scholars with limited network data to apply the model's insights into the interaction of media and networks to their empirical cases. Figure 1 provides a visualization of the four network types I use: the Small World, the Village (or Clique), the Opinion Leader, and the Hierarchical network. All ties in these networks are assumed symmetric: anyone you influence also influences you.¹³ Network structure is set at the time internal motivations are distributed and is assumed constant as participation spreads.¹⁴ However, I explicitly vary network structure in the analysis. Unless stated otherwise, the distribution of net internal motivations is uncorrelated with the position of individuals in the network.

The Small World network (Watts 1999) captures life in cities and suburbs. Individuals' local networks overlap significantly, so that people connected to each other are likely to share other connections as well. Yet each person also may have a number of "weak ties" (Granovetter 1973), defined here as connections that link socially distant people. This gives individuals a chance to influence people outside their own clusters, allowing for the swift spreading of information across the network. These networks can form when tight-knit groups of friends or family disperse upon, for example, transition to college or forced migration.

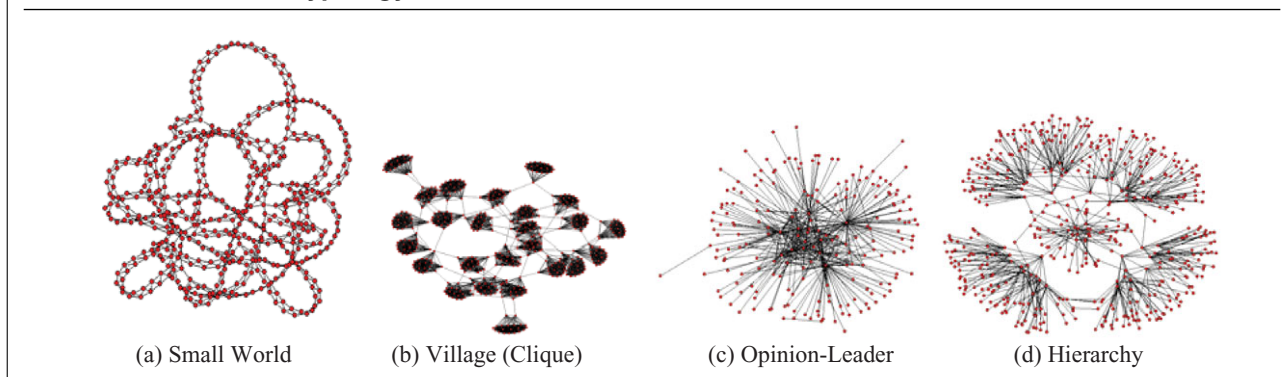
The Village network is meant to mimic villages, small towns, and cliques, in which everyone knows everyone else within the social unit. In addition, individuals may have weak ties to people in other villages; local elders

¹² While I could have presented results for the weak class instead, its lower mean participation levels make interactions less easy to display on the page.

¹³ Note that symmetric ties do not imply symmetric influence within the network. One opinion leader directly affects the decisions of all of her followers at once, while significant numbers of these followers must act in concert to induce her to change her behavior.

¹⁴ I assume that social networks are comparatively stable and exist independently of their role in affecting one instance of political participation or opinion formation. In other words, the rate of network change is slow compared to the spread of participation.

¹¹ Siegel (2011a, 2011b) discuss variant models in which individuals cease participating in response to repression.

FIGURE 1. Network Typology

who regularly meet in multivillage conclaves might possess these. Such individuals are assumed to be few, and so paths of influence are strongly clustered within social units.

In these two networks individuals have roughly equal connectivity, and thus may be thought of as possessing approximately equal levels of social influence. Each of these networks has two parameters: one dictating average connectivity—the average number of people to whom one is connected—and one dictating the frequency of weak ties. I find that the effect of the media in these nonelite networks depends on, in order of importance: (1) average connectivity, and (2) whether the number of weak ties is suboptimal, optimal, or greater than optimal in encouraging participation.¹⁵ I thus vary both connectivity and the optimality of weak ties in subsequent analysis.

I also consider two elite networks in which individuals have unequal connectivity. In the first, the Opinion Leader network, most people have one or two connections, while a few have many. Simple versions of this network have also been termed “stars” or “wheels” (e.g., Gould 1993). Networks of this form have both elites and followers. The elites have many more connections than the followers and so have much more influence. This sort of network might be premised on expertise of the elites, or arise naturally from individuals’ preferring to befriend already popular people. In addition to the case in which internal motivations are distributed uncorrelated with number of connections, for these networks I also analyze cases in which elites have uniformly high (positive correlation) or uniformly low (negative correlation) internal motivations. I refer to both correlations as the case of unified elites.

The Hierarchy is the second elite network, and is a standard organizational structure that is common in bureaucracies as well as many social and religious groups. Hierarchies expand outward in width from the top, with each individual connected to one superior and a set number of subordinates. Individuals also may be connected to others in their same organizational

level to capture bonds formed from close work. The power of elites within the Hierarchy lies in their privileged placement at its top, rather than in their absolute number of connections; a too-wide (or too-narrow) Hierarchy can actually diminish the influence of the elites. I analyze cases in which elites have uniformly high (positive correlation) or uniformly low (negative correlation) internal motivations, as well as cases in which internal motivations are uncorrelated with position in the network.

The Hierarchy has two parameters, while the Opinion Leader network has one. One parameter in both networks translates into the level of influence of the elites, which is a function of the number of elites and the structure of their connections. The second parameter in the Hierarchy translates into the level of influence of the followers, which is significant when there are many within-level connections. I find that the effect of the media in these elite networks depends on, in order of importance: (1) the correlation of motivations with network positions, (2) the level of influence of their elites, and (3) the level of influence of their followers (only for the Hierarchy). I thus vary all three in subsequent analysis.

Media

As I discuss the model of the media at length in the previous section, here I only connect the verbal description there with the details provided in this section. A media outlet is any source of information that exists outside of the social network, is accessible by all within the network, and influences individuals’ external motivations. Each media outlet is defined by two parameters: its strength, M_S , and its bias, M_B . There may be more than one outlet, and each one may have different values of these parameters.

The baseline media (comprising any number of unbiased outlets) provide information about the global participation rate, i.e., the aggregate rate of participation across the population.¹⁶ As noted above, the strength of the media, M_S , determines how much each individual

¹⁵ Optimal is defined functionally as the parameterization that yields the highest level of participation in the intermediate class, all else equal. This is often not maximal (Siegel 2009).

¹⁶ Or, equivalently, individuals accurately infer the global participation rate from the baseline media’s fair reporting.

weights the media's information relative to that arising from one's social network. I analyze the baseline media in the next section by exploring the simultaneous variation of media strength, network properties, and motivation class.

Biased media act in a simple manner: they add their bias, M_B , to the aggregate participation rate and provide that information instead.¹⁷ This biased information is used by individuals in the same manner as is unbiased information. Both types of bias, polling bias and advocacy, are implemented identically. The only difference is that I keep polling bias relatively small, whereas media advocates for or against participation have bias sufficient to set their information to 1 or 0, respectively, for all global participation rates.

In the fifth section I analyze two cases of biased media. First I consider one biased media outlet, and then two biased media outlets on opposite sides of the issue. I call L the one biased toward participation and R the one biased toward the status quo. To accommodate this I add another parameter, M_L , which dictates the relative strength of the proparticipation L outlet. The net information from two biased media outlets is a linear combination of the biased outlets' information, with M_L the weight on L's information. The net information is combined with one's local information as for an unbiased outlet.

In my analysis I simultaneously vary media strength, network properties, media bias, and, for two outlets, the strength of the L outlet. Though I keep my analysis to two biased outlets, it can easily be extended to multiple biased outlets with the addition of parameters dictating their relative strengths.

This analysis assumes that the media have identical strength for all people, regardless of their personal characteristics. This assumption is not likely to hold when media are obviously biased. Consequently, in the sixth section I allow for selection into biased media and repeat my analysis under this new assumption. How selection works depends on the number of media outlets. With one biased media outlet, those with high (low) internal motivations raise (lower) their individual values of media strength in the presence of a proparticipation (antiparticipation) media outlet. With two media outlets, those with high (low) internal motivations increase (decrease) the weight on L's information and decrease (increase) the weight on R's information.

Analysis relies on simulation to overcome the problem of intractability. In this section I have detailed the progression of each realization of the model, from the distribution of internal motivations and creation of the network to the attainment of the steady state. However, as the realizations of both individual motivations and network connections involve random components, a single realization is insufficient. Thus, for each set of parameters I compute 200 realizations and take the mean of all steady-state outcomes. I report mean aggregate participation in the next three sections. The seventh

section summarizes all results of the model, discusses measurement issues, and derives and discusses related testable hypotheses. Readers may find it useful to read that section before the three prior results sections, particularly if their focus is on substantive applications of the model.

THE MEDIA: SPUR AND SPOILER

Media Strength—Uncorrelated Networks

In this section I consider baseline media that provide unbiased information equal to the global rate of participation. I vary media strength: the level of influence the media have relative to social networks. I first consider uncorrelated networks: nonelite networks and elite networks in which elites are not unified.

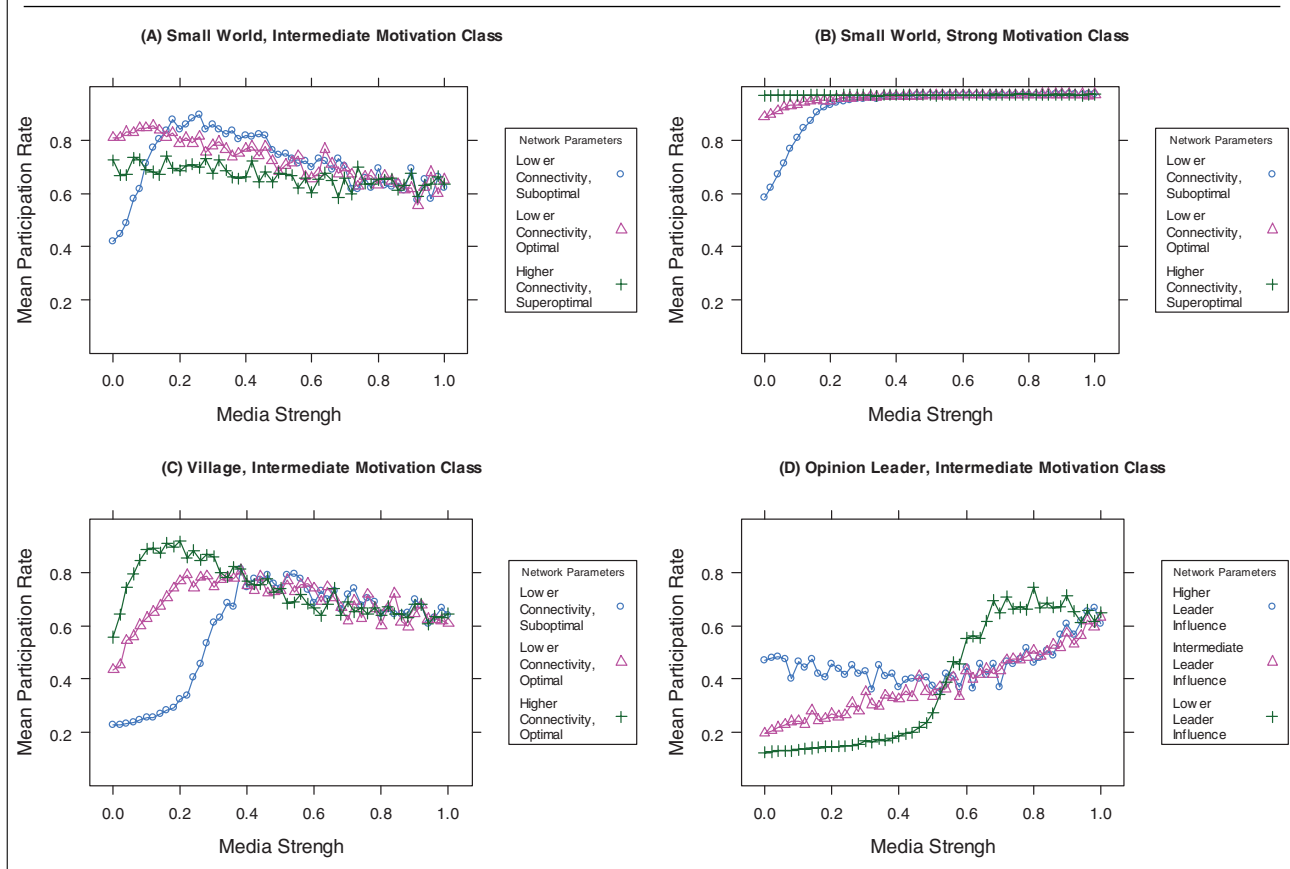
Figure 2 displays four plots of the mean equilibrium participation rate against the strength of the media for uncorrelated networks. Each plot illustrates a different combination of network type and motivation class. Lines within each plot show representative values of network parameters. First examine Figures 2A and 2B, which correspond to participation within Small World networks in the intermediate and strong motivation classes, respectively. One difference is striking. In Figure 2A, increasing the strength of the media has either a negative or a nonmonotonic effect on equilibrium participation. In Figure 2B, increasing media strength solely increases participation.

To understand this difference, consider the role played by unbiased media. By making known the global participation rate, the media present to individuals a fuller picture of the population. Since people want to participate more if their fellows do, this more complete picture can have different effects depending on both network structure and the distribution of motivations in the population (motivation class). In the strong motivation class there are more highly motivated people, and hence more early participants. The media make these individuals known to those who might otherwise only know about local nonparticipants; thus the media act as a spur. This turns out to be a robust regularity. For all network types, in all network parameterizations, increasing the strength of the media in the strong class increases aggregate participation. Only the rate and amount of increase vary with network type and parameterization. Further, the rate of increase in the strong class is generally faster in the same circumstances in which increasing media strength in the weak or intermediate classes most increases participation. Given this, we will consider only the intermediate class in what follows.¹⁸

In contrast, in the intermediate class (and in the weak as well), the media's role is far more conditional on network structure. The optimal level of media strength—that which maximizes participation—is no longer always at $M_S = 1$; now increasing the strength

¹⁷ The sum of the global participation rate and bias cannot be less than 0 or exceed 1. Alternatively, individuals might infer an incorrect global participation rate from biased media reporting.

¹⁸ Corresponding figures for the strong class to all those shown below are available on request.

FIGURE 2. Motivations and Optimal Media Strength in Uncorrelated Networks

of the media can diminish participation over a significant range of media strength. Compare Figures 2A, 2C, and 2D, corresponding to Small World, Village, and Opinion Leader networks, respectively, all in the intermediate class and with uncorrelated motivations. For the Small World network only when connectivity is low and weak ties suboptimal can the media have any positive effect on participation, and even here only when the media are less than one-fourth as influential as social networks. In the Village network the media increase participation for some strengths, but begin to be a hindrance well before they are as influential as social networks. In the Opinion Leader network the media are more helpful, particularly when leaders are not as influential.

Within the intermediate class, aggregate participation depends on an often delicate trade-off between encouraging enclaves of participation and allowing for the spread of participation out of these enclaves (Siegel 2009). In the absence of the media, only network neighbors matter. An individual whose neighbors are highly internally motivated may be spurred to participate, even though few apart from these neighbors are participating. Enclaves such as this encourage participation, which can then spread to others. Media that reach all individuals heavily favor participatory spread at the expense of local encouragement. Strong media can inform individuals with participating neighbors about others' lack of participation. Suddenly participation

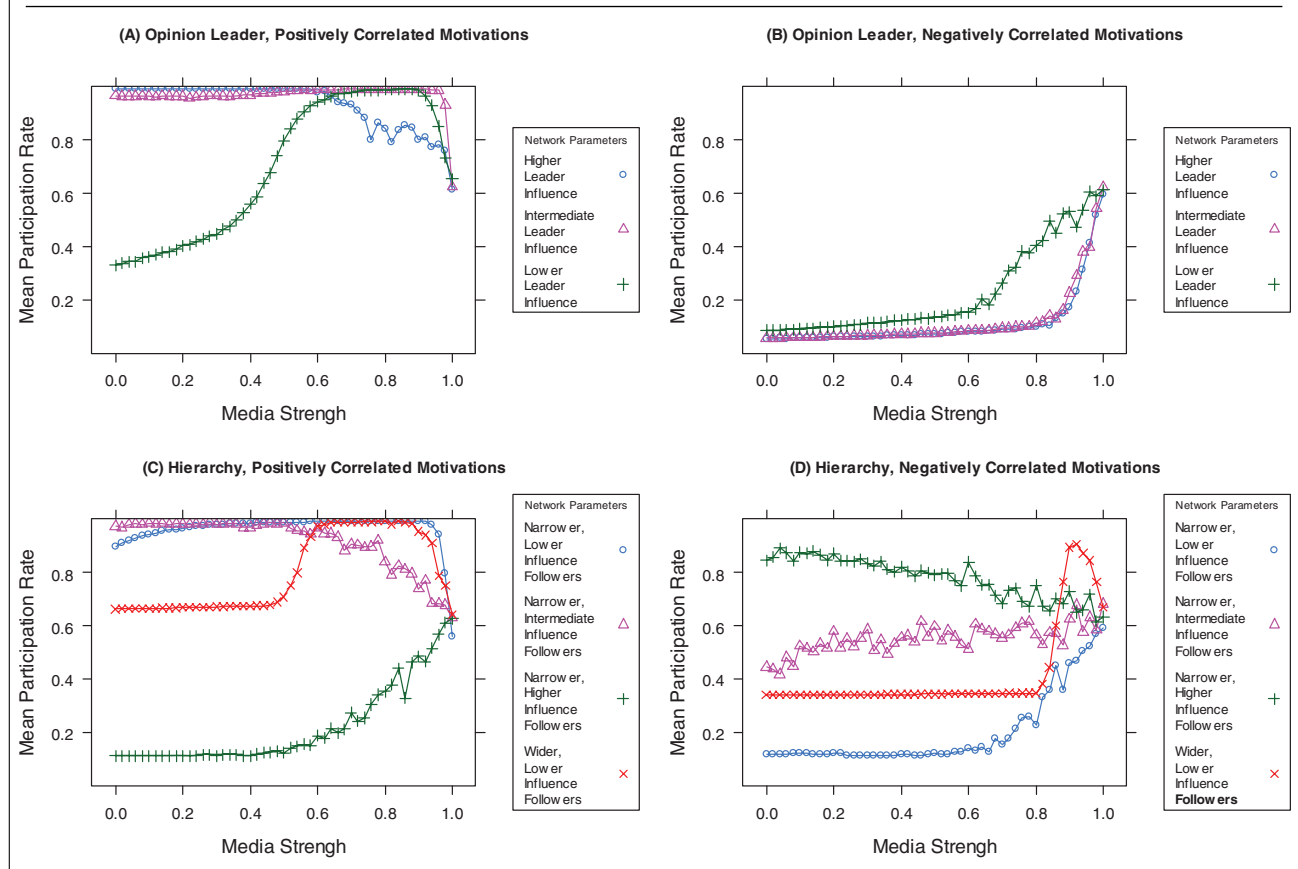
doesn't look as good to those individuals, and the enclave doesn't form. In the aggregate, decisions like this can result in a spoiler role for the media, in which too much media influence hinders the encouraging role of the social network.

The question, then, is when do the media spur participation, and when do they act as spoilers? In general, the answer is conditional on network type and structure. Thus, in the aggregate, we cannot simply say that the media substitute for or complement social networks, and cannot aggregate individual-level influences on behavior. However, some generalities hold. Specifically, the poorer a particular network configuration is at spreading behavior, the more helpful the media are, and the more media strength can be increased before it starts to diminish participation. Uncorrelated Opinion Leader networks typically are less good at spreading participation than uncorrelated Hierarchies or Villages, which are less good than Small World networks. All else equal, then, we should expect the media's ability to spur participation to decline by network type in this same pattern.

Media Strength—Unified Elite Networks

Thus far we have discussed the media's influence in the context of networks in which elites are not unified in support for an option. Figure 3 extends these

FIGURE 3. Optimal Media Strength in Unified Elite Networks



results to unified elite networks. Figures 3A and 3B show equilibrium participation rates as a function of media strength when elites in Opinion Leader networks are positively and negatively correlated, respectively. Figures 3C and 3D display the same for Hierarchies. As in uncorrelated networks, we see that non-monotonic responses to media strength are the norm here. Only in the negatively correlated Opinion Leader network are the media purely a spur to participation.

There is a difference, though. For nearly all parameterizations of unified elite networks, the level of media strength necessary for the media to have a noticeable effect on aggregate participation is higher than in uncorrelated networks. In many cases the media do not begin to play a role until they are equally as influential as social networks, and in some cases the media must be *fifty times* more influential than social networks before they even begin to alter participation levels. In such cases, elites unified in encouraging mass participation are able to drive aggregate behavior as long as they receive a modicum of attention from their followers. If the media are less influential in individual decision-making than social networks (Beck et al. 2002), then the media will rarely have any effect at all when elites are unified in support for either option.

One can see why this occurs by considering a simple “star” Opinion Leader network, in which the central opinion leader is participating and there are many followers. Each follower’s neighborhood only consists of

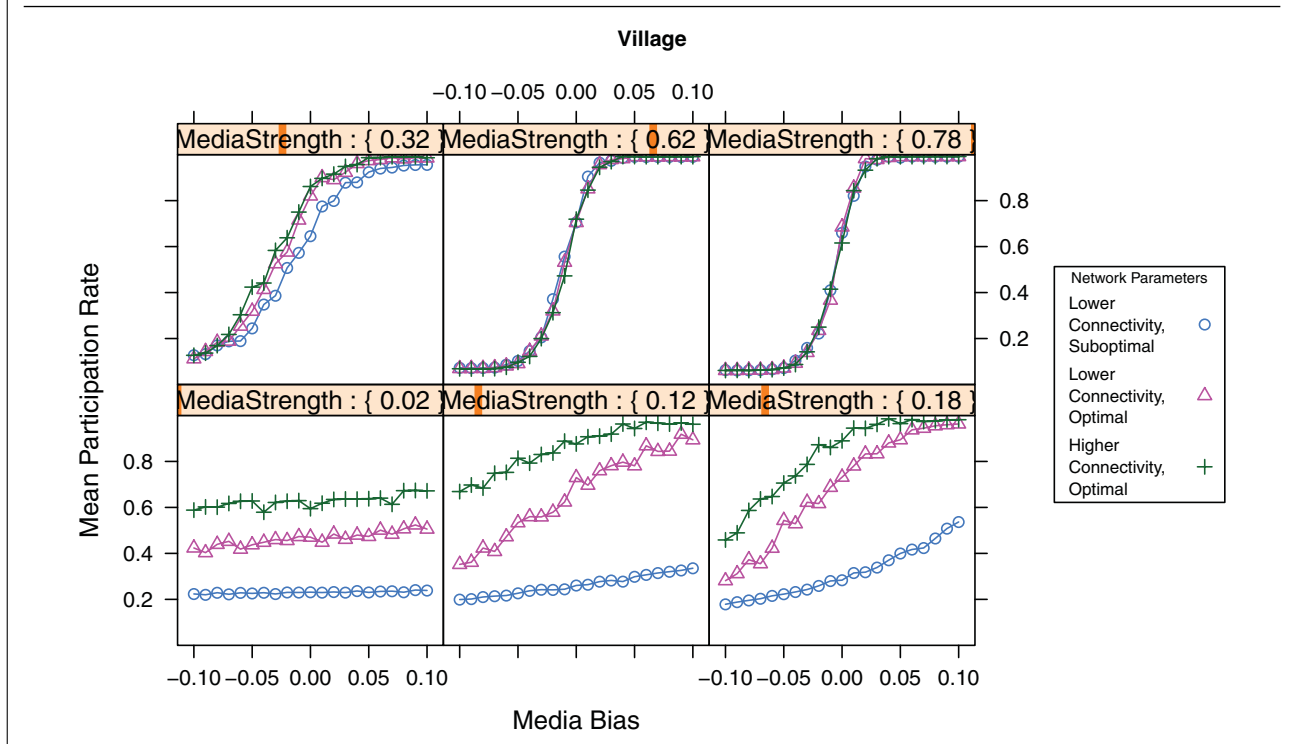
a single participating elite in this case. In this case, even media information weighted nine times more than that from this elite still results in external motivations great enough to allow for the immediate participation of a substantial range of individuals.¹⁹ From these early movers, participation quickly spreads. Unified elites send a similarly unified signal to their followers, echoing this simple star example.

This result should not be understated. The power of unified elites in networks is such that an external source of accurate information is unlikely to play a role in mass behavior, calling into question the utility of external propaganda within these types of societies. If leaders are sufficiently influential and all desire some behavior on the part of their followers, there is virtually nothing a realistic media outlet can do by passing along accurate information about the population, even if the outlet is trustworthy.

In fact, not only can the media seemingly do little to stop participation unified elites from getting their way, it can unwittingly *aid* them when the structure of the network prevents the elites from reaching the entire population with their influence. When opinion leaders are too few or a Hierarchy is too wide, the population can become fragmented. Clusters split from the

¹⁹ Recall that one participates if the sum of external and internal motivations exceeds a cutoff. Higher external motivations allow the cutoff to be exceeded with a wider range of internal motivations.

FIGURE 4. Polling Bias in Nonelite Networks



leadership are insufficiently tied to participants, resulting in their low external motivations. In such cases the media encourage participation by effectively connecting the clusters. The media and networks act in synergy here: first the people in power get a subset of the population to participate, and then news of this subset is passed on to those not so favorably connected, spurring them on. This cycle continues until participation is near total. The truth, in this case, is effectively manipulated by elites to overcome limitations in their network of influence.

BIASED MEDIA

One Media Outlet—Polling Bias

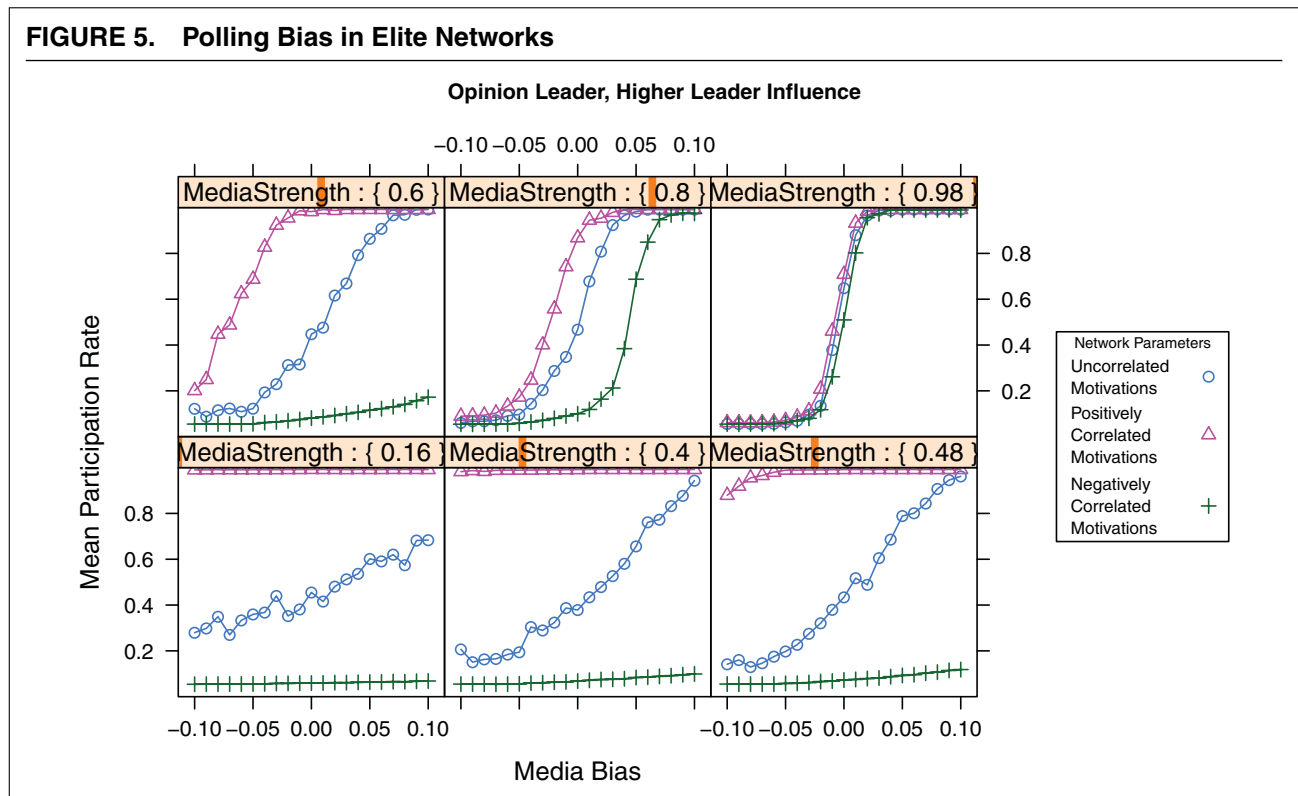
I begin exploring the effect of biased media by looking at polling bias with a single media outlet. Here the media’s information is biased by some relatively small amount. Figures 4 and 5 each display six plots of participation against media bias, conditional on media strength. (The bar running along the top of each plot indicates the level of media strength for that plot.) Though I will focus on bias in the range of $\pm 3\%$, the horizontal axes take bias out to $\pm 10\%$. In all cases, bias acts as expected: positive bias increases participation and negative bias decreases it. For more interesting results, consider Figure 4 first, which illustrates participation in Village networks. The upper-left plot uses $M_S = 0.32$, for a media outlet roughly half as influential as one’s network, within the range suggested by data (Beck et al. 2002). At this strength the lines correspond-

ing to different network parameterizations almost converge, and do converge at even higher strengths. Consequently, sufficiently strong biased media render most differences between network parameterizations irrelevant for predicting participation.

The magnitude of the effect of bias is striking at this level of strength. Consider the lower connectivity, suboptimal weak tie network parameterization. With no bias, its mean equilibrium participation rate is 65%. With each additional 1% of negative bias this drops to 57%, then 51%, then 39%. Even low levels of polling bias can thus substantially decrease participation. By the same token, a positive bias of only 3% takes participation to 88%. Polling bias is therefore a potent determinant of aggregate outcomes, and a substantial threat to the proper functioning of democratic governments. For example, a swing in bias from -3% to $+3\%$ (in this noncompetitive media environment) alters aggregate participation (e.g., turnout or support for one of two candidates) by 49%, enough to swing most elections. Because the level of bias remains small, the media outlet can occasion this swing without necessarily lying to the public: error is within $\pm 3\%$, as stated in the fine print.

The power of media bias stems from its action early in the spread of participation, when aggregate participation is still small. Positive bias subtly shifts the playing field, increasing the number of individuals willing to participate early. These early participants encourage others via both social networks and the media, increasing aggregate participation. Then, this new level of global participation is again inflated, leading to yet more participation, and so on until a cascade of

FIGURE 5. Polling Bias in Elite Networks



participation occurs. Negative bias acts in reverse, hiding the activities of the highly motivated and so preventing the initiation of any cascade. The plots in the bottom row of Figure 4 indicate that even weaker media outlets can effectively use bias, by virtue of its role in path-dependent behavior (Page 2006).²⁰

What can limit the damage arising from biased media? One answer, explored below, is competing media. A second answer, shown in Figure 5, is unified elites. Figure 5 mirrors Figure 4, except in an Opinion Leader network. The central line displays the uncorrelated case under high leader influence as a baseline. We see again that media bias is powerful in uncorrelated networks: a +3% bias increases participation from 43% to 61% when $M_S = 0.48$, as seen in the bottom right plot. Look, however, at the other two lines, which correspond to a negatively and a positively correlated Opinion Leader network, again under high leader influence. At this same level of media strength, bias only begins to have an effect when it is large, and even then the effect is comparatively small. Only when the media outlet has more influence than one’s own social network do lower levels of bias begin to diminish the power of elites. Unified social elites are thus a solution to a biased media outlet, though in a sense

one is merely replacing one monolithic influence for another.

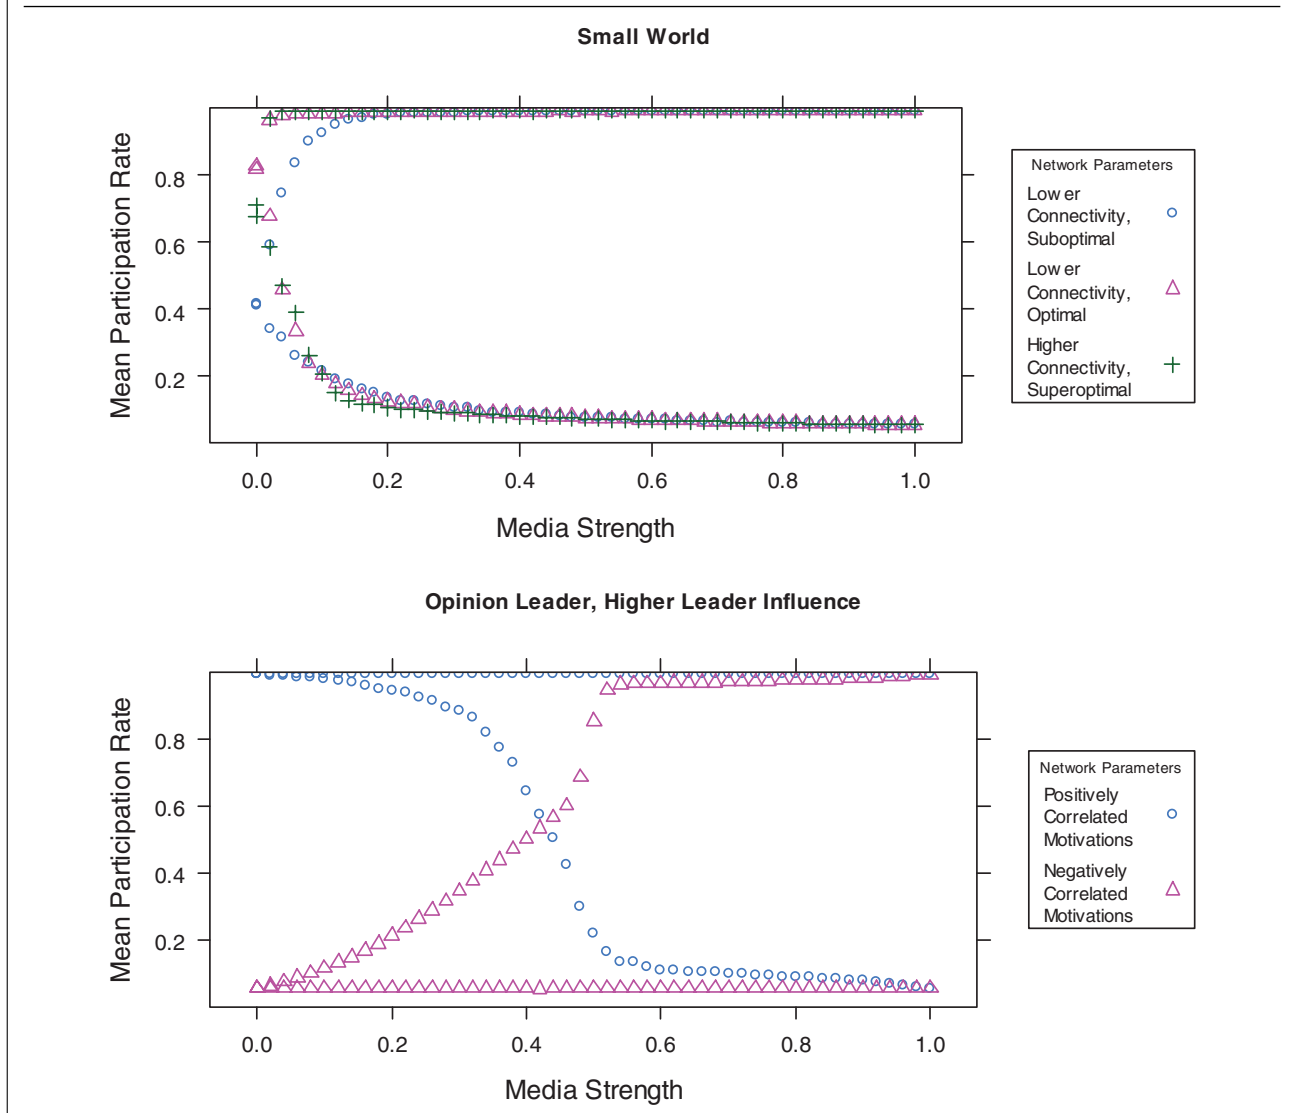
One Media Outlet—Advocacy

Polling bias can occur under the radar, given that it is a slightly altered informative message. This might be desirable for biased media outlets looking to maintain widespread trust. Advocacy, however, is apparent to all. Media outlets such as the old partisan papers, *Fox News*, or MSNBC, or many Internet news sources openly advocate for a particular option and so do not hide their biases. Positively biased advocates come out for participation, while negatively biased advocates argue for the status quo. Figure 6 displays two plots of participation against media strength. The top plot is of three parameterizations of a Small World Network, the bottom of the positively and negatively correlated cases of an Opinion Leader network under high leader influence. Each line corresponds to one network parameterization; the upper (lower) portion of a line represents positive (negative) bias.

The top plot illustrates the incredible effectiveness of advocacy. Even low-strength media—which may be all one can expect given open bias—is sufficient to drive participation much higher or much lower, dominating any role of network parameterization. This is again due to the path-dependent nature of the process. The manner in which the media’s information affects external motivations, and, therefore, the decision to participate, implies that advocacy functions in much

²⁰ This also offers a general caution to be wary of potential bias in all fundamentally path-dependent behavior. Actions such as bias that influence early events in path-dependent behavior have their impact effectively multiplied over time, so that even minor bias can lead to substantial shifts in outcome.

FIGURE 6. Advocacy (upper (lower) Portion of Each Line Reflects Positive (negative) bias)



the same way as uniformly increasing (or decreasing) the internal motivations of the population. This has a strong positive (or negative) effect on participation levels, since it significantly affects the number of early movers whose behavior is translated into increased (or decreased) aggregate participation in the network.

Two differences distinguish advocacy from polling bias: (1) Its greater bias implies that it alters participation levels more strongly and requires less media strength to be effective. (2) As seen in the bottom plot, an influential media advocate is able to partially overcome the power of social elites. While an advocate in the presence of unified elites requires more strength for its influence to be effective than it does in the absence of unified elites, advocacy still plays a role at a level of media strength that makes media less influential than social networks. Consequently, even social elites are not proof against an openly biased media outlet, as long as it is sufficiently influential. Given this, I now turn to the other means by which a

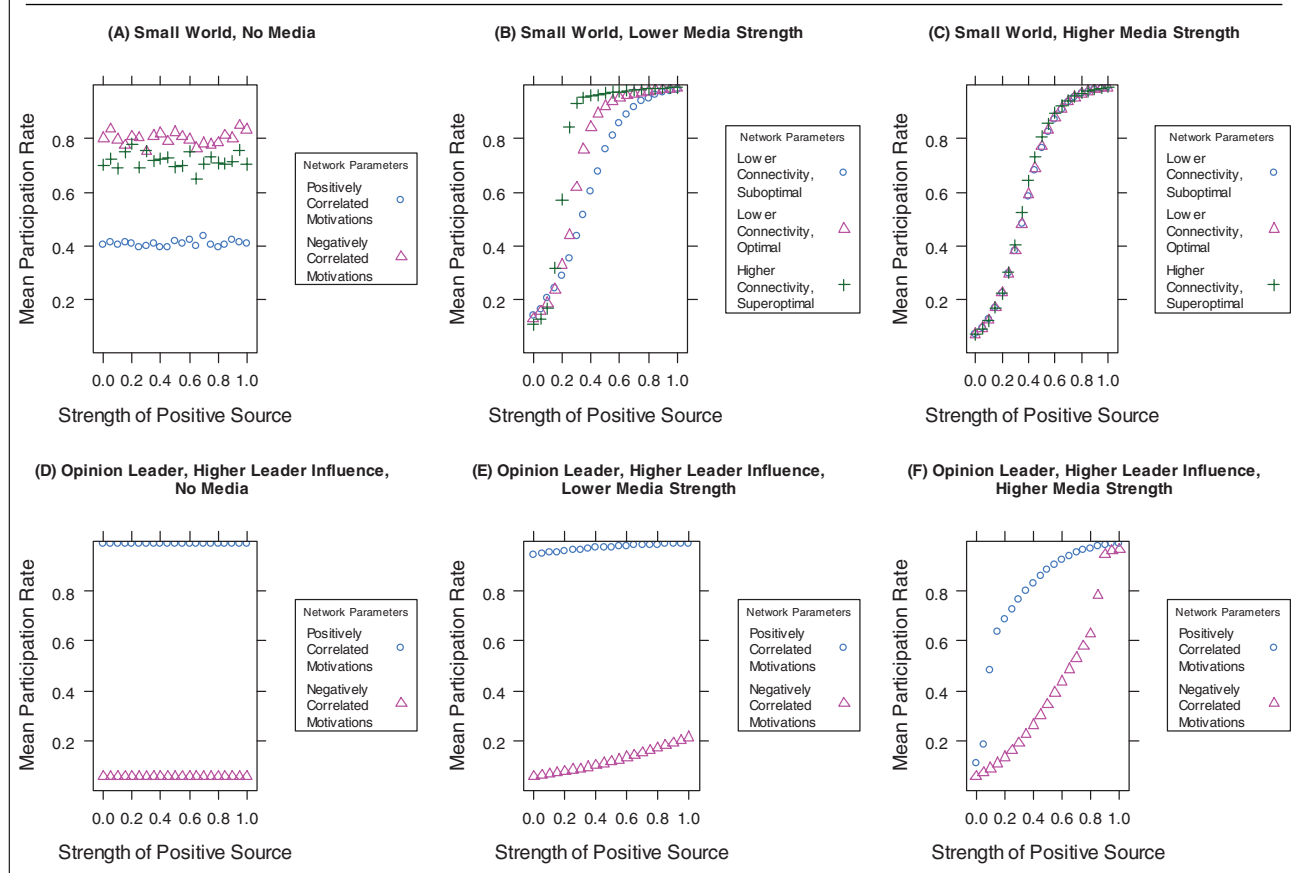
biased media outlet might be neutralized: countervailing media.

Multiple Media Outlets

We have seen that bias is powerful. So powerful, in fact, that one would expect to observe additional biased media outlets whenever there were no institutional constraints prohibiting them. Competition for attention will often lead them to take opposite positions, and here I explore one such case. I assume two oppositely biased outlets, L and R, with L (R) positively (negatively) biased.²¹ However, the analysis extends in a straightforward fashion to more than two outlets of varied bias, at the cost of additional parameters representing the relative influence of the outlets.

²¹ Recall that the information from each outlet is added via a linear combination to yield the net media information, with M_L the weight on L's information.

FIGURE 7. Advocacy—Two Opposing Media Sources



Consider polling bias first. Here each media outlet sends information equal to the global participation rate plus its bias. This implies that the net information from the media is equal to the global participation rate plus a linear combination of each outlet’s bias. This linear combination is the net polling bias.²² This is equivalent to a single biased media outlet employing the net polling bias. Given this, the results described in and around Figures 4 and 5 already apply to the case of multiple media sources employing polling bias.²³ For outlets of equal strength, if L’s bias exceeds R’s the net media information drives participation higher; if R’s exceeds L’s the net information drives participation lower. If we assume that biased sources will seek the highest level of bias they are able to slip by their audiences and that individuals tolerate the same level of bias from each outlet, the media outlets’ biases will cancel out, providing accurate net information. However, if L (R) had greater strength than R (L), or L (R) could maintain a similar level of strength with greater bias than could R (L), then the net message would drive

participation higher (lower). In other words, the media outlet better able to maintain its strength despite bias has the advantage.

Thus, multiple media outlets can, under some conditions, be a panacea for polling bias. What about advocacy? Because advocates effectively ignore the global participation rate, the net media information is now just a linear combination of 1 (full participation) and 0 (no participation). Since the relative strength, M_L , multiplies the former, the net information is just M_L .²⁴ In general, the relative strength of the proparticipation outlet will not be equal to the global participation rate, so two advocates—even equally balanced ones—will *not* see their influences cancel out. Instead, they are equivalent to one outlet sending intermediate information between 0 and 1.

Figure 7 illustrates the effect of intermediate information. As usual, the vertical axis reports equilibrium participation rates; the horizontal axis now displays the relative strength of L (M_L). As in Figure 6, the top plots illustrate participation given three representative parameterizations of a Small World network, while the bottom illustrate participation given elites with positively and negatively correlated motivations. The middle two plots (B and E) look at a lower level of media

²² To see this, let the global participation rate be p_i and $M_B(L)$ and $M_B(R)$ be the respective biases of L and R. Then the net media information, S , is $S = M_L * (p_i + M_B(L)) + (1 - M_L) * (p_i + M_B(R)) = p_i + [M_L * M_B(L) + (1 - M_L) * M_B(R)]$. The term in the square brackets is the net polling bias.

²³ Additional media outlets would only increase the number of terms in this linear combination, not change the analysis.

²⁴ $S = M_L * 1 + (1 - M_L) * 0 = M_L$. More than one advocate on each side would not change the analysis.

strength ($M_S = 0.2$), the top-right and bottom-right plots (C and F) look at a higher level of media strength ($M_S = 0.6$), and the top-left and bottom-left plots (A and D) provide results without media ($M_S = 0$) as a baseline.

The most important commonality in the plots is that there is an asymmetry between the positive and the negative outlet. In all cases, the net effect of two opposing outlets with equal strength ($M_L = 0.5$) is an *increase* in participation over the baseline. It is easier, and so requires a weaker media outlet, to spur participation (or to drive support away from the status quo) than it is to quell participation (or maintain the status quo). In a democracy, one would therefore expect it to be easier for a biased media to agitate contrary to the incumbent government than in support of it, coloring what media strategies would be effective by the incumbent government and the positioning of successful media outlets.

The reason for this asymmetry is that intermediate information has two effects. First, early in the spread of participation, when global participation is still low, intermediate information encourages more people to participate. Second, when participation is widespread, it discourages people from participating. The first effect is more powerful than the second because of the path-dependent nature of the process: the first ensures that there will be more participants in each person's social network early in the process, allowing network interactions to spread participation. Once participation is widespread any discouragement will have a limited effect, since it also has to overcome behavior in one's social network. There will be similar, though more involved, logic when we discuss selection into media in the next section.

There is a more subtle thing to note as well in Figure 7. Without unified elites intermediate information is effective; however, unified elites can more easily limit the effect of intermediate information than they could the stronger bias of a single advocate. Thus, unified elites facing influential and biased media would do well to encourage the growth of a countervailing media outlet. However, as we see in Figure 7F, were the media to gain in strength, two equally strong media outlets would diminish the ability of elites to dictate aggregate outcomes, so elites must be cautious against this development.

SELECTION INTO MEDIA

In the preceding analysis all individuals treated the media as identically influential, despite the possibility of bias. This obviously was a simplification: the media's influence might vary by substantive issue in a way that social networks' influence, often arising from the networks' greater roles in everyday life, does not. Further, this variation may depend on individuals' preferences: individuals may be more likely to trust media sources that support their proclivities. Accordingly, it is important to understand the effects of *selection* into media influence. The flexibility of the basic model allows the direct incorporation of selection. In this section I ana-

lyze a simple model of selection, and use it to illustrate both the robustness of earlier results to selection and the way selection alters the effect of media bias on aggregate behavior.

To incorporate selection in the model, I assume that an individual's internal motivations dictate the degree to which that person is influenced by a biased media outlet. The more someone's internal motivations exceed the mean in the population, the more that individual is influenced by media outlets with positive (participation) biases, and the less that individual is influenced by media outlets with negative (status quo) biases. The same is true in reverse for individuals with motivations less than the population mean. With one media outlet individuals change the strength they assign to the media (M_S) to accommodate changes in influence.²⁵ With two media outlets individuals change the strength they assign to outlet L (M_L), while keeping the overall influence they assign to the media (M_S) the same. In other words, under selection individuals pay diminished attention to any outlet that doesn't match their preferences, and pay more attention to any outlet that does. Technical details are provided in the Online Appendix.

Analysis of selection comprises a repetition of the analysis in the previous section, but with selection incorporated in the model. There is one straightforward conclusion of this analysis: the model is robust to the addition of selection. Figures A1–A4 in the Online Appendix display corresponding results to Figures 4–7; their similarities are apparent. To derive more interesting conclusions, however, it is necessary to understand the complicated causal mechanisms that drive the effect of selection. This is the focus of the rest of this section. To improve clarity, I discuss below whether selection increases participation, decreases participation, increases bias, or decreases bias, leaving off the important final clause: *relative to the case of media bias without selection*. In other words, all discussion of selection's effects that follow are relative to identical scenarios lacking only selection.

Selection into media has two opposing effects that take slightly different forms depending on the number of media outlets. First consider two outlets. On the one hand, selection increases the net media information observed by those with high internal motivations. These individuals assign more weight to the positively biased outlet and less to the negatively biased outlet. In the aggregate, this results in higher levels of participation: more such people are led by the greater media information to participate early in the process. On the other

²⁵ One could also imagine time-varying selection into an unbiased, informative media outlet, based on the particular message the outlet is promulgating at that moment and one's own predisposition to believe it. Though I do not consider it here, one can reason through its outcomes. Highly motivated individuals would select into ignoring the media early on, since it would report little participation; this would drive up early participation as compared to an equivalent unbiased media outlet without selection. Once aggregate participation was high individuals with low motivation would begin to ignore it. However, by then these individuals' social networks would, on average, provide the same information of widespread participation. Thus, this sort of selection should increase aggregate participation.

TABLE 1. Effect of Selection into Media

	No Leaders	Unified Leaders
Polling Bias, 1 Outlet	↑ bias	↓ bias (at high str)
Advocacy, 1 Outlet	↑ bias	↑ bias (at low str), ↓ bias (at high str)
Polling Bias, 2 Outlets	↑ participation	↓ bias (at high str)
Advocacy, 2 Outlets	↑ participation	↑ participation

Note: Cell values of “↑ (↓) bias” indicate that selection increases (decreases) the effect of bias (i.e., leads to increased (decreased) participation under positive bias and decreased (increased) participation under negative bias), as compared to the case without selection. Modifiers “(at high str)” and “(at low str)” indicate that the corresponding effect of selection only occurs when media strength is relatively high and relatively low, respectively. Cell values of “↑ participation” indicate that selection weakly increases participation for both directions of bias.

hand, selection decreases the net media information observed by those with low internal motivations. These individuals assign less weight to the positively biased outlet and more to the negatively biased outlet. In the aggregate, this results in lower levels of participation as people later in the process of behavioral spread now fail to be mobilized due to the smaller media information.

The net of these two effects determines the aggregate outcome under selection, and varies according to the structure of the network and the type of media bias. In general aggregate dependence on selection can be complex; however, previous analyses can provide some expectations. In particular, the first, positive effect of selection acts earlier in the process than the second, negative effect of selection. Consequently, the demonstrated importance of early participation suggests that the positive effect of selection should dominate, absent specific network properties that render the negative effect particularly potent. This suggests increased participation under selection in general.

To check this suggestion’s accuracy one can compare aggregate participation under selection to that in the absence of selection; I do this in Figures A5– A10 of the Online Appendix. All plots display aggregate participation under selection *minus* aggregate participation without selection; positive values thus indicate that selection increases participation. Two conclusions of this analysis stand out. First, selection has a positive effect more often than not, and almost uniformly when there are two media outlets. Two, selection’s effect is not always positive, indicating further complexity in selection arising from network properties. I discuss this complexity briefly here, after first summarizing selection’s effect in Table 1.

Each cell in Table 1 details how the addition of selection alters aggregate behavior given biased media. The last two rows, corresponding to the case of two media outlets, bear out the above logic. In almost all cases, selection increases participation for the reasons given above. In the one case in which it does not—two media outlets, unified leaders, and polling bias—selection instead diminishes the effect of bias. This only happens, however, at high media strength, and only when media bias works oppositely to leaders’ unified preferences. The cause of this deviation from

increased participation is the structure of the elite network. Under selection elites with positively correlated motivations pay less attention to negative bias, driving participation higher and in the opposite direction of the bias. Participation can be driven lower too, at very high media strengths (when the media are four or more times as influential as social networks). This occurs when elites with negatively correlated motivations experience positive bias. Here selection implies that elites see smaller net media information. This reduces the effect that media bias has on the borderline members of the elite (those with relatively higher, though still low internal motivations) who might otherwise be spurred to participate by the combination of media bias and participatory followers. Less responsiveness by borderline elites leads to less participation overall, reducing the effect of bias. In essence, the effect of selection on elites’ responses to the media matters more than its effect on followers’ responses, leading to a decrease in the impact of bias.²⁶

Now consider one media outlet. Again there are two opposing effects of selection, for the same reasons as with two media outlets. However, because selection operates directly on individuals’ media strengths when there is only one media outlet, the magnitude of each of the opposing effects is smaller.²⁷ Since the positive effect drives results more than the negative in general, decreasing both effects tends to diminish the net positive effect of selection, relative to what we see with two media outlets.

The first two rows of Table 1 illustrate this. Without elites, the overall effect of bias is enhanced. Selection still induces more participation in the presence

²⁶ This does not show up in the case of advocacy due to the more potent effect of advocacy at high media strength, which overwhelms elites’ impact. This is similar to that observed without selection in Figure 7.

²⁷ Net media information, which is what selection effectively changes when there are two media outlets, only affects external motivations in one way: it alters the term multiplying media strength. In contrast, selection with one media outlet directly alters media strength, which appears in that same term but also in the $(1 - M_S)$ term multiplying the local participation rate. Changing media strength moves these two terms in opposite directions, which accounts for the smaller effect of selection with one media source.

of positive bias; however, selection now induces less participation in the presence of negative bias. With the net positive effect of selection diminished, avenues open up for bottlenecks in network structure to lead to diminished participation. The relative importance of specific connections is what is important: worse connectivity and/or suboptimal numbers of weak ties imply more potential bottlenecks, and the more bottlenecks, the more selection can enhance negative bias.²⁸ This is because network bottlenecks and the negative effect of selection both act to inhibit behavioral spread in a complementary fashion. Bottlenecks limit the avenues across which behavior can spread, and selection makes it less likely that behavior will spread across any existing avenue.

With unified leaders and one media outlet the story is very similar to that observed in polling bias with two media outlets. Specifically, at high media strengths, where the media actually can begin to overcome social elites' influence, selection results in less consideration of the opposing biased media outlet by elites. This reduces the effect of bias. The only new twist here is that, under advocacy at low media strengths only, selection can enhance bias. Advocacy begins to counteract elites' influence starting at low media strengths by spurring followers on (positive bias) or driving followers away (negative bias). Selection can aid this when the advocate is insufficiently strong to do it on its own; however, once it is, elites again become the focal point of selection as noted.

In sum, though selection into media does not alter the main results on media bias presented in the previous section, it does have some interesting, if complex, interactions with network structure that the model can tease out. Selection can increase participation or enhance or diminish bias, depending on the type of bias, the nature of the media, and the structure of the network, highlighting again the importance of considering network structure. Substantively, for most cases in which there is more than one media outlet we again see the advantaged state of participation relative to the status quo. Thus selection can exacerbate some of the issues that arise under biased media.

EMPIRICAL IMPLICATIONS AND SUMMARY OF RESULTS

This article presents an abstract model with broad substantive consequences relating to the interaction between social network structure and media influence and bias. Four things are necessary to derive and test empirical implications of the model. First, one must make sure the model applies. This requires that two primary assumptions hold: (1) There must be a dichotomous set of options over which individuals are

²⁸ This occurs when media strength is high enough to otherwise eliminate network differences. The top row of Figure A5 in the Online Appendix supports this for Village networks. Small World networks, with fewer bottlenecks, display this pattern more weakly. The pattern almost completely inverts under polling bias for a network in which everyone is connected (and so lacks bottlenecks).

TABLE 2. Heuristics for Network Type Measurement

Network Type	Identification Heuristics
Small World (SW)	Most Know Someone Far Away, Info Travels Quickly
Village/Clique (V)	Clumped Social Groups, Poor Info Travel btwn Groups
Opinion Leader (OL)	Few Elites Drive Opinion, Info from Common Elites
Hierarchy (H)	Defined Organizational Structure with Elites at Top

choosing. One of these must be a status quo option, and everyone in a population must begin supporting the status quo. (2) Support for the non-status-quo option, which I have termed *participation*, must be increasing in others' support. As noted previously, participation includes, but is not limited to, support for a new policy alternative, political participation in voting or social movements or protests, or voting for a challenger.

Second, one must identify the network type, as this helps both to constrain the range of applicable hypotheses and suggest what additional data might be needed. Table 2 provides a few heuristics for identifying each network type.

Third, one must derive a hypothesis as to the effect of media influence and/or bias on aggregate participation, conditional on network structure and individuals' proclivities to participate. These hypotheses will connect a measure of the aggregate rate of participation (the dependent variable) to one or more measures of the model's parameters (the independent variables). The model's analysis, performed over the previous three sections, structure these hypotheses. Table 3 summarizes this analysis and provides a sense of the range of hypotheses the model generates. Specifically, the table specifies the way in which the model's nonmedia parameters condition (1) the effect on participation of increasing media strength (the second column), and (2) the efficacy of media bias in altering participation (the third column). Note that because the model's implications are robust to the presence of selection, this table does not include selection. Table 1 in the previous section contains a summary of the effect of selection.

Fourth and finally, one must determine which data are necessary to test one's hypothesis and measure these data. One of the virtues of the model is that its data requirements can often be comparatively weak compared to the complexity encapsulated in the model. In particular, it is not equally important to gather data on all parameters in all cases. Table 4 provides two potentially helpful pieces of information in identifying the most useful data. The middle and bottom rows of the center column respectively specify the conditions under which increasing media strength is most likely to increase participation and under which media bias is most effective. This highlights for which values of the model's parameters one would expect to see the most

TABLE 3. Parameters' Effects on Media's Impact on Participation

Parameter (Effect on:)	Increasing Media Strength	Efficacy of Media Bias
Motivation Class	Greater Participation	Minor Effect
Average Connectivity	Lesser Participation More Likely	Minor Positive Increase
Weak Ties	Lesser Participation More Likely	Minor Positive Increase
Elite Influence	Lesser Participation More Likely (Varies for Correlated Hierarchy)	Varies Nonmonotonically
Follower Influence (H)	Lesser Participation More Likely (Varies for Correlated Hierarchy)	Increase (Generally)
Elite Unity	Generally Minor Effect	Large Decrease

Note: Increasing Motivation Class equates to moving toward the strong class; increasing Elite Unity means moving to positively or negatively correlated motivations. The middle column specifies the effect of an increase in the parameter in the left column on the degree to which increasing media strength increases participation. The rightmost column specifies the effect of a parameter increase on the efficacy of media bias.

TABLE 4 Effect of Media Parameters

Information	When Increasing Most Effective	Most Important Data
Media Strength	Uncorrelated Networks, Strong Motivations, Poorly Connected Networks	Elite Unity, Motivation Class, Network Type, Network Parameters
Media Bias	Uncorrelated Networks, Better-Connected Nonelites	Media Strength, Elite Unity, Network Type, Network Parameters (OL,H only)

effect of the media on participation. The rightmost column ranks the most important data to gather for understanding the media parameters' effects, given the information in the center column. Once one has chosen which to data collect it becomes necessary to measure the model's parameters. Section 4 in the Online Appendix provides guidelines for the measurement of all model parameters and a summary table for reference.

To make this procedure more concrete and better connect theory to empirics, I illustrate how one may apply each of the four steps to derive and test the three hypotheses discussed in the Introduction, though I do not conduct empirical tests here. I immediately dispense with the first step as all three hypotheses were chosen due to their relevance to the model, and treat the remaining three steps in the context of each hypothesis.

First, will the media be more effective in caucus states than in primary states, assuming caucus states encourage less communication across precincts? This assumption suggests that the relevant network in caucus states is the Village, while the relevant network in primary states is the Small World. This is step 2. Table 3 suggests one conditional hypothesis in its center column: outside of the strong motivation class, increasing media strength will be less likely to increase participation the better connected the nonelite network is. Specifically, the model shows that uncorrelated Opinion Leader networks typically are less good at spreading participation than uncorrelated Hierarchies or Villages, which

are less good than Small World networks. Thus, all else equal, we should expect the media's ability to spur participation to decline by network type in this same pattern. Consequently, as long as states are not in the strong motivation class, then the media should be more effective in caucus states than in primary states. This is step 3.

To test this hypothesis one must determine the necessary data and measure it, as in step 4. The center cell in the rightmost column of Table 4 provides us a list; for nonelite networks, data in order of importance are motivation class, network type, and network parameters. Measuring the first and the third requires survey instruments, as described in Section 4 of the Appendix; we have already identified the second. However, all we really want to know is if the population of either group of states is in the strong motivation class. Indicators such as significant early participation may be sufficient to suggest the strong class. Conversely, if early participation is not significant, or if the population is very small (see the Appendix), we might feel comfortable not measuring motivation class at all. Regardless, we do not need to measure network parameters to test this hypothesis; nor must we measure media strength. In short, only two pieces of data are needed to test this hypothesis: a measure of aggregate participation and the type of system in the state.

Second, when can external media outlets affect the behavior of insular elite networks? Aggregate opinion in insular groups can differ markedly from that in

the larger population, and even the latter's unbiased media can appear biased to those in the network. We have already specified that we are considering elite networks; this further argument suggests that we are interested in hypotheses related to bias. Looking at the third column of Table 3, we see that the most important nonmedia parameter is elite unity. Specifically, the model indicates that change is likely to happen only when elites in the network are not unified. However, analysis also indicates that all elites need not be on board with change, and that when elites are divided the media can play a substantial role by taking advantage of network connections. This is our hypothesis. To test it, we note that the bottom-right cell in Table 4 suggests that the two most important pieces of data are media strength and elite unity, and these are sufficient to test our hypothesis.²⁹ The former may be measured via analysis of latent variables in a battery of survey instruments getting at trust in and attention to media outlets and social networks, the latter via public statements and elite surveys.

Third, how will decreased state ability to limit access to external messages from media and NGOs alter authoritarian states' strategies in maintaining power? Here the model suggests two strategies, depending on the type of network. In an elite network, the state must ensure unified elites matching its status-quo bias. In a nonelite network, the state must minimize connectivity within and across disparate groups. Both claims can be drawn from the center column in Table 3 and are summarized in the center cell of Table 4.³⁰ Testing these claims requires data on network types and either elite unity for elite networks or average connectivity and weak ties for nonelite networks, as well as data on state strategies. The network data for nonelite networks can be obtained via surveys measuring discussant number and the degree of commonality of discussants for a sample of the population. As I noted in the Introduction, China's behavior appears in line with these claims (King, Pan, and Roberts 2012).

CONCLUSION

Social networks mediate our interactions with all aspects of the political sphere. We ask those to whom we are connected what they think about topics of public import and whether information is accurate or relevant; we consider their behavior as potential models for our own; we worry about their responses to our own behavior and opinions. We have seen that this not only implies that social networks and the mass media both affect individual behavior, but also that *social networks condition media's impact in the aggregate*. One often can't determine the effect of increased influence of an honest, informative media outlet without knowing the structure of the social network in place. Social networks

can also amplify media bias, leading to large swings in aggregate behavior made more severe when individuals can select into media outlets. Countervailing media outlets and unified social elites with common preferences can mitigate the effect of bias; however, media outlets promulgating anti-status-quo bias have an advantage.

These conclusions derive from a model based on very general assumptions, making the conclusions relevant to a broad array of substantive settings. The model's applicability to any dichotomous action with a status quo implies that the model speaks to questions in public opinion, vote choice, and voter turnout, as well as participation in social movements, mass protests, and insurgencies. Further, since a media outlet in the model can correspond to a conventional media outlet such as *The New York Times* or *FOX News*, or to state propaganda, NGOs, national party leadership, Internet personalities, and the like, one can formulate and test hypotheses designed to answer questions such as those that opened the Introduction. I worked through three such hypotheses in detail in the previous section, but there are many more I could have discussed. For example, media pundits can drive mass opinion even when individuals can select their media outlets; they are not constrained to preach to the choir.

Such broad applicability means that the model's theoretical conclusions also have substantial empirical consequences. Most notably, failing to account for the conditioning role of social networks implies potential bias in empirical studies of media effects in the aggregate. Additional data collection on networks is necessary to overcome this bias; this also helps in forecasting the effects of changes in the media. The model's mixture of qualitative and quantitative network specifications will aid in this endeavor.

Finally, a natural extension of the model would be to include media as strategic players who choose bias to increase their share of the market (Mullainathan and Shleifer 2005). While I don't analyze this here, the existing model suggests some likely outcomes and new questions that can be addressed. For example, how will increased ability to select from numerous media outlets affect aggregate public opinion? The model indicates that selection tends to enhance anti-status-quo bias, implying an increasingly oppositional public opinion. This is likely to be further served in any media market by additional oppositional media outlets. This might fragment the media's message, but it is also possible that the proliferation of media outlets could lead to increased media influence, since at least one of the outlets might effectively cater to one's predilections. This would result in either strong opposition to the status quo that would need to be managed by the state, or homophily of networks as pro-status-quo people separate themselves from the influence of those with divergent preferences. The latter connects endogenous network formation and the creation of new media outlets. It could form the start of a more complete model of opinion formation that fully accepts the demonstrated interdependency of social networks and mass media.

²⁹ In nonelite networks, testing the efficacy of bias requires measurement of only one parameter, media strength.

³⁰ I am assuming strategies that do not involve changing individuals' fundamental preferences, and thus am not considering motivation class.

Supplementary materials

To view supplementary material for this article, please visit <http://dx.doi.org/10.1017/S0003055413000452>

REFERENCES

- Ahn, T.K., Robert Huckfeldt, and John Barry Ryan. 2010. "Communication, Influence, and Informational Asymmetries among Voters." *Political Psychology* 31 (5): 763–87.
- Althaus, Scott L. 2003. *Collective Preferences in Democratic Politics: Opinion Surveys and the Will of the People*. Cambridge, UK: Cambridge University Press.
- Ansolabehere, Stephen, Roy Behr, and Shanto Iyengar. 1991. "Mass Media and Elections: An Overview." *American Politics Quarterly* 19 (1): 109–39.
- Barabas, J., and J. Jerit. 2009. "Estimating the Causal Effects of Media Coverage on Policy-Specific Knowledge." *American Journal of Political Science* 53 (1): 73–89.
- Bartels, Larry M. 1996. "Uninformed Votes: Information Effects in Presidential Elections." *American Journal of Political Science* 40 (1): 194–230.
- Beaudoin, C.E., and E. Thorson. 2004. "Social Capital in Rural and Urban Communities: Testing Differences in Media Effects and Models." *Journalism & Mass Communication Quarterly* 81 (2): 378–99.
- Beck, P.A. 2002. "Encouraging Political Defection: The Role of Personal Discussion Networks in Partisan Desertions to the Opposition Party and Perot Votes in 1992." *Political Behavior* 24 (4): 309–37.
- Beck, P.A., R.J. Dalton, S. Greene, and R. Huckfeldt. 2002. "The Social Calculus of Voting: Interpersonal, Media, and Organizational Influences on Presidential Choices." *American Political Science Review* 96 (01): 57–73.
- Berelson, Bernard R., Paul F. Lazarsfeld, and William N. McPhee. 1954. *Voting: A Study of Opinion Formation in a Presidential Election*. Chicago: University of Chicago Press.
- Carmines, Edward G., and James A. Stimson. 1990. *Issue Evolution: Race and the Transformation of American Politics*. Princeton: Princeton University Press.
- Centola, Damon, and Michael Macy. 2007. "Complex Contagions and the Weakness of Long Ties." *American Journal of Sociology* 113 (3): 702–34.
- Chiang, Chun Fang, and Brian Knight. 2008. "Media Bias and Influence: Evidence from Newspaper Endorsements." Working Paper.
- Chong, Dennis. 1991. *Collective Action and the Civil Rights Movement*. Chicago: The University of Chicago Press.
- Cohen, Jeffrey E. 2004. "If the News Is So Bad, Why Are Presidential Polls So High? Presidents, the News Media, and the Mass Public in an Era of New Media." *Presidential Studies Quarterly* 34 (3): 493–516.
- DellaVigna, S., and E. Kaplan. 2007. "The Fox News Effect: Media Bias and Voting." *The Quarterly Journal of Economics* 122 (3): 1187–234.
- Delli Carpini, Michael X., and Scott Keeter. 1996. *What Americans Know About Politics and Why it Matters*. New Haven, CT: Yale University Press.
- Druckman, James N., and Kjersten R. Nelson. 2003. "Framing and Deliberation: How Citizens' Conversations Limit Elite Influence." *American Journal of Political Science* 47 (4): 729–45.
- Druckman, James N., and Michael Parkin. 2005. "The Impact of Media Bias: How Editorial Slant Affects Voters." *Journal of Politics* 67 (4): 1030–49.
- Druckman, J.N. 2004. "Political Preference Formation: Competition, Deliberation, and the (ir) Relevance of Framing Effects." *American Political Science Review* 98 (4): 671–86.
- Eveland Jr, W.P., A.F. Hayes, D.V. Shah, and N. Kwak. 2005. "Understanding the Relationship between Communication and Political Knowledge: A Model Comparison Approach using Panel Data." *Political Communication* 22 (4): 423–46.
- Gentzkow, Matthew. 2006. "Television and Voter Turnout." *The Quarterly Journal of Economics* 121 (3): 931–72.
- George, Lisa M., and Joel Waldfogel. 2006. "The New York Times and the Market for Local Newspapers." *American Economic Review* 96 (1): 435–47.
- Gerber, Alan, Dean S. Karlan, and Daniel Bergan. 2009. "Does the Media Matter? A Field Experiment Measuring the Effect of Newspapers on Voting Behavior and Political Opinions." *American Economic Journal: Applied Economics* 1 (2): 35–52.
- Gerber, Alan S., Donald P. Green, and Christopher W. Larimer. 2008. "Social Pressure and Voter Turnout: Evidence from a Large-scale Field Experiment." *American Political Science Review* 102 (1): 33–48.
- Gould, Roger V. 1993. "Collective Action and Network Structure." *American Sociological Review* 58 (2): 182–96.
- Graber, Doris A. 2004. "Mediated Politics and Citizenship in the Twenty-First Century." *Annual Review of Psychology* 55: 545–71.
- Granovetter, Mark S. 1973. "The Strength of Weak Ties." *American Journal of Sociology* 78 (6): 1360–80.
- Groeling, Tim, and Samuel Kernell. 1998. "Is Network News Coverage of the President Biased?" *Journal of Politics* 60: 1063–87.
- Huckfeldt, Robert. 2001. "The Social Communication of Political Expertise." *American Journal of Political Science* 45 (2): 425–39.
- Huckfeldt, Robert R., and John Sprague. 1995. *Citizens, Politics and Social Communication: Information and Influence in an Election Campaign*. Cambridge, UK: Cambridge University Press.
- Jerit, Jennifer, Jason Barabas, and Toby Bolsen. 2006. "Citizens, Knowledge, and the Information Environment." *American Journal of Political Science* 50 (2): 266–82.
- Kahn, Kim Fridkin, and Patrick J. Kenney. 2002. "The Slant of the News: How Editorial Endorsements Influence Campaign Coverage and Citizens' Views of Candidates." *American Political Science Review* 96 (2): 381–94.
- King, Gary, Jennifer Pan, and Margaret Roberts. 2012. "How Censorship in China Allows Government Criticism but Silences Collective Expression." Tech. rep., Working Paper, 2012. copy at <http://j.mp/LdVXqN> Export Tagged XML BibTex Google Scholar.
- Klofstad, Casey. 2011. *Civic Talk: Peers, Politics, and the Future of Democracy*. Philadelphia: Temple University Press.
- Klofstad, Casey A., Anand E. Sokhey, and Scott D. McClurg. 2013. "Disagreeing about Disagreement: How Conflict in Social Networks Affects Political Behavior." *American Journal of Political Science* 57 (1): 120–34.
- Kuran, Timur. 1991. "Now Out of Never—The Element of Surprise in the East European Revolution of 1989." *World Politics* 44 (1): 7–48.
- Ladd, Jonathan McDonald, and Gabriel S. Lenz. 2008. "Exploiting a Rate Shift in Communication Flows to Document News Media Persuasion: The 1997 United Kingdom General Election." Working Paper.
- Lake, Ronald La Due, and Robert Huckfeldt. 1998. "Social Capital, Social Networks, and Political Participation." *Political Psychology* 19 (3): 567–84.
- Lau, Richard R., and David P. Redlawsk. 1997. "Voting Correctly." *American Political Science Review* 91 (3): 585–98.
- Leighley, Jan E. 1990. "Social Interaction and Contextual Influences On Political Participation." *American Politics Research* 18 (4): 459–75.
- Lupia, Arthur and Mathew D. McCubbins. 1998. *The Democratic Dilemma: Can Citizens Learn what they Need to Know?* Cambridge, UK: Cambridge University Press.
- McAdam, Doug. 1986. "Recruitment to High-Risk Activism—The Case of Freedom Summer." *American Journal of Sociology* 92 (1): 64–90.
- McClurg, Scott D. 2003. "Social Networks and Political Participation: The Role of Social Interaction in Explaining Political Participation." *Political Research Quarterly* 56 (4): 449–64.
- McClurg, Scott D. 2006. "The Electoral Relevance of Political Talk: Examining Disagreement and Expertise Effects in Social Networks on Political Participation." *American Journal of Political Science* 50 (3): 737–54.
- McLeod, J.M., K. Daily, Z. Guo, W.P. Eveland Jr, J. Bayer, S. Yang, and H. Wang. 1996. "Community Integration, Local Media Use, and Democratic Processes." *Communication Research* 23 (2): 179–209.

- Mondak, Jeffrey J. 1995a. "Media Exposure and Political Discussion in US Elections." *Journal of Politics* 57 (1): 62–85.
- Mondak, Jeffrey J. 1995b. *Nothing to Read: Newspapers and Elections in a Social Experiment*. Ann Arbor, MI: University of Michigan Press.
- Mullainathan, Sendhil, and Andrei Shleifer. 2005. "The Market for News." *American Economic Review* 95 (4): 1031–53.
- Mutz, D.C., and P.S. Martin. 2001. "Facilitating Communication Across Lines of Political Difference: The Role of Mass Media." *American Political Science Review* 95 (1): 97–114.
- Mutz, Diana C. 2002. "The Consequences of Cross-Cutting Networks for Political Participation." *American Journal of Political Science* 46 (4): 838–55.
- Niven, David. 2002. *Tilt?: The Search for Media Bias*. Westport, CT: Praeger Publishers.
- Page, Benjamin I. 1996. *Who Deliberates?: Mass Media in Modern Democracy*. Cambridge, UK: Cambridge University Press.
- Page, Scott E. 2006. "Path Dependence." *Quarterly Journal of Political Science* 1 (1): 87–115.
- Petersen, Roger Dale. 2001. *Resistance and Rebellion: Lessons from Eastern Europe*. Studies in rationality and social change, Cambridge, UK: Cambridge University Press.
- Rojas, H., D.V. Shah, and L.A. Friedland. 2011. "A Communicative Approach to Social Capital." *Journal of Communication* 61: 689–712.
- Rolfe, Meredith. 2012. *Voter Turnout: A Social Theory of Political Participation*. Political Economy of Institutions and Decisions. Cambridge, UK: Cambridge University Press.
- Ryan, John B. 2011. "Social Networks as a Shortcut to Correct Voting." *American Journal of Political Science* 55 (4): 753–66.
- Siegel, David A. 2009. "Social Networks and Collective Action." *American Journal of Political Science* 53 (1): 122–38.
- Siegel, David A. 2011a. "Non-Disruptive Tactics of Suppression Are Superior in Countering Terrorism, Insurgency, and Financial Panics." *PLoS ONE* 6 (4): e18545.
- Siegel, David A. 2011b. "When Does Repression Work? Collective Action in Social Networks." *Journal of Politics* 73 (1): 232–47.
- Sinclair, Betsy. 2012. *The Social Citizen*. Chicago: University of Chicago Press.
- Sokhey, Anand E., and Scott D. McClurg. 2012. "Social Networks and Correct Voting." *The Journal of Politics* 74 (3): 751–64.
- Watts, Duncan J. 1999. *Small Worlds: The Dynamics of Networks between Order and Randomness*. Princeton Studies in Complexity. Princeton: Princeton University Press.
- Zaller, John R. 1992. *The Nature and Origins of Mass Opinion*. Cambridge, UK: Cambridge University Press.