Online Appendix for *Social Networks and the Mass Media*

David A. Siegel*

This appendix contains six sections. The first section details and provides further justification for the parameters of the model, including selection into media. It also specifies the order of operations within the model. The second details the method of network construction and parameterization. The third lays out and justifies the methodology by which results given in the paper were derived. The fourth section provides guidelines for measuring the parameters in order to empirically test hypotheses drawn from the model. The fifth section offers additional results for media bias under selection. The sixth section provides results for a model variant in which there is an additional parameter, media penetration, which dictates whether or not someone has access to the media. In other words, it relaxes the assumption that the media's information is seen by all.

1 Model Parameters and Order of Operations

The model includes two types of actors: a finite population of N individuals connected via a social network, and one or more mass media outlets external to these networks. To keep the focus on the role of the media and its interaction with network structure, this model generalizes the model of Siegel (2009) through the addition of the media. The model is analytically intractable even without the media, necessitating the use of computational methods to analyze it. Computational models allow for the inclusion of numerous outcome measures and input parameters. However, results from complex models can be difficult to interpret, potentially diminishing the insight one can derive from them (Axelrod 1997). To minimize this problem, I keep the model of the media simple, using only three parameters, and focus on a single dependent variable: the steady-state (equilibrium) aggregate rate of participation in the population. As noted in the text, participation refers to support for the option (of a dichotomous pair of options) that is not the status quo. All results given in the text or this appendix are in terms of aggregate participation.

Individual Behavior

Each of the N people in the population has incentives to participate separated into two disjoint components. The first, called *internal motivation*, includes all factors relating to one's desire to participate that do not depend on the participation of others. Examples of

*Associate Professor, Department of Political Science, Duke University, Durham, NC 27708; e-mail: david.siegel@duke.edu; web: `http://people.duke.edu/~das76`.

factors that might fall into this category include a moral certainty in the value of participation, prior beliefs about the benefits of one of the options, or an opportunity cost derived from participating or choosing a particular option. Internal motivation will most likely be the sum of random variables drawn from many different distributions, each corresponding to a different aspect of one's motivation. Assuming these draws are independent, and that their third central moments satisfy the Lyapunov condition, the sum of these random variables will be distributed normally, by Lyapunov's central limit theorem. Let each person i's internal motivation be b_i, and the mean and standard deviation of the normal distribution from which they are drawn be respectively b_{mean} and b_{stdev}.

Note that, even if one argues that there are few relevant motivations or that draws are not independent, this only affects the specification of the distribution from which the b_i are drawn.[1] The basic idea of each individual's having some net internal motivation—upon which this model rests—still holds. The b_i may be uncorrelated with network position, positively correlated (large b_i get assigned to the most well-connected people in the Opinion Leader network or to those at the top of the Hierarchy), or negatively correlated (vice versa). See Siegel (2009) for additional details and justification of internal motivations.

The second component of behavior is one's *external motivation*, denoted $c_{i,t}$ for each individual i, at each time t. It includes all factors relating to one's desire to participate that depend on the participation of others. The model assumes a parsimonious form for $c_{i,t}$ in the absence of the media: $c_{i,t+1} = lpr_{i,t}$, with $lpr_{i,t} \in [0, 1]$ the local participation rate, which is the participation rate of those to whom one is directly connected, or equivalently of those within one's (immediate) social network. External motivations start at 0 for all people as all people begin not participating, i.e., $c_{i,0} = 0$ for all i.[2] This ensures that external motivations are increasing in the participation of others within one's social network.

In the presence of the media one incorporates information from both social networks and the media according to the following equation: $c_{i,t+1} = (1 - (M_S)_i)lpr_{i,t} + (M_S)_i * S_i$, where $(M_S)_i$ is the media strength parameter for person i and S_i is the net information the media transmits to person i. Media strength measures the relative influence of the media as compared to one's social network. The subscript on each parameter indicates that under selection media influence and net information may vary by individual. Absent selection, all individuals experience the same media strength (i.e., $(M_S)_i = M_S$) and net information (i.e., $S_i = S$).

The media's net information varies by the number and bias of media outlets. A general equation for one or two outlets is: $S_i = (M_L)_i * S_L + (1 - (M_L)_i) * S_R$. Here L and R correspond to the pro-participation and pro-status-quo outlets, respectively, S_L and S_R are the information released by each outlet, and $(M_L)_i$ is the relative influence of the L outlet. The last is subscripted again due to selection and is identical for all individuals in the absence of selection (i.e., $(M_L)_i = M_L$). Let p_t be the global (i.e., aggregate) participation rate. Then the information released by each outlet is: $S_j = \max\{0, \min\{1, p_t + (M_B)_j\}\}$ for $j = \{L, R\}$,

[1]Different distributions do lead to different aggregate levels of support (see, e.g., Granovetter and Soong 1986; Yin 1998) in the absence of the media; an investigation of the same with the media could be a productive avenue of further research.

[2]Initial external motivations are set at their minimum to avoid hard-wiring participation into the model. The maximum value of $c_{i,t}$ is 1.

where $(M_B)_j$ is the bias of media outlet j. I typically assume bias is positive for outlet L and negative for outlet R. Though each can vary continuously, in my analysis I assume $(M_B)_L = -(M_B)_R = M_B$ for reasons given in the text. A single media source is represented by setting $M_L = 1$ (or $M_L = 0$, equivalently), and an unbiased media is represented by setting $(M_B)_j = 0$. The maximum and the minimum in the equation for S_j ensure that the information stays between 0 and 1. In polling bias I keep the level of bias low so that global participation still usually affects the information. In advocacy I make it high enough so that global participation never affects the information.

Under selection the assignment of the parameters $(M_S)_i$ and $(M_L)_i$ is not identical for all individuals. Rather, it varies according to the relative internal motivation of the individual in the following fashion. First, relative motivation is defined by $r_i = \frac{b_i - b_{mean}}{3b_{stdev}}$. When $M_L = 1$ (or $M_L = 0$) so there is one media source, one's relative motivation can either be consonant with (both positive or both negative) or not consonant with (one positive and the other negative) the bias. In the former case, M_S increases the further internal motivations are from the mean according to: $(M_S)_i = M_S * \max\{0, (1 - |r_i|)\} + \min\{1, |r_i|\}$. In the latter case M_S decreases the further internal motivations are from the mean according to: $(M_S)_i = M_S * \max\{0, (1 - |r_i|)\}$. When $M_L \neq 0$ so there are two media sources, one can be more supportive of the positive (L) or negative (R) source and less supportive of the other. This is accomplished by simultaneously increasing the term corresponding to the more preferred source (M_L or $1 - M_L$) and decreasing the one corresponding to the less preferred source ($1 - M_L$ or M_L) in the same manner the strength of the single source was increased or decreased. Note that the sum of the strengths still adds to one.

Each period, every individual decides whether or not to participate in that period. The decision rule is simple: A person i participates at time t if and only if $b_i + c_{i,t} > 1$. That is, one participates if and only if her net motivation to do so exceeds 1.[3] As the left-hand side of this inequality is increasing in others' participation, this rule implies that the more people who participate, the more one wants to do so as well.

As the b_i are unbounded while $c_{i,t}$ is bounded by 0 and 1, there can always be rabble-rousing types (Granovetter 1978) who participate regardless of their fellows. All early participation arises from those with $b_i > 1$. On the other hand, those with $b_i \leq 0$ will never participate under any circumstances. I bound $c_{i,t}$ so that internal motivation, which encompasses personal characteristics, is never completely washed out, and that there are always individuals who may never want to participate and who always want to participate.

Mass Media

The mass media comprise some number of media outlets. A mass media outlet is an external source of information that, along with information about the participation of those within one's social network, affects one's net external motivation. The media is not a part of the network; rather, media outlets send information to all individuals regardless of network position. Thus, any type of network may be served by any combination of media outlets, and media advocates and social network elites may have different preferences over options.

[3]The cutoff of 1 is arbitrary but unimportant for analysis. Different individual values of internal motivation set different thresholds for each individual, and the mean of the distribution of these can be varied to overcome different cutoffs in an identical manner. In other words, it is unnecessary to vary both the cutoff and the mean of the distribution of internal motivations.

Though it is not a strategic actor within the model, intent can be read into the choice of bias by the media. Two characteristics (and three parameters) define the media: *strength*, (M_S, M_L) and *bias*, (M_B).

The media's strength, M_S, is a proxy for the relative level of trust and attention the media commands compared to one's social network. The strength of the pro-participation source, M_L, is a proxy for the relative level of trust and attention the pro-participation media outlet commands compared to the pro-status-quo outlet. I follow the literature in representing the media's role by using the global rate of participation as a proxy for an unbiased, informative media's message (Oliver and Myers 2003).[4] In other words, an unbiased, informative media sends a message to all that effectively transmits the participation rate across the whole population. The form of the external motivation indicates that setting the strength of the media to 1 obviates social networks, as local participation is no longer a factor, while a media strength of 0 obviates the media.[5] I assume that the strength of the media is the same for all members of the population when there is no selection into media.[6] As strength is essential for the media's functioning, in my analysis I first vary strength on its own.

Media bias has already been discussed in the context of individual behavior. The media is informative, and the baseline media's unbiased message is equal to the global rate of participation, p_t. As discussed in the text, this may be thought of as an unbiased poll of population support for the non-status quo option or a proxy for fair reporting on an issue. Here "fair" is defined not as truthful in an objective sense, but as a fair picture of the information that is in the population as a whole about an issue. If this information is flawed, the media's report may lack accuracy, but if it is not then fair reporting may approximate accurate reporting.

Positive (negative) bias suggests a global level of participation higher (lower) than the true value and, via the form of $c_{i,t}$, drives individuals' external motivations higher (lower) than they would have been given an unbiased media. Note that polling bias is not dichotomous; biased media outlets can attempt to add (or subtract) larger or smaller amounts from the true global participation rate. In order to mitigate concerns that individuals will adjust their level of trust in or attention to a biased media, leading to a correlation between bias and strength, I examine only low levels of polling bias—on the order of normal polling error—which we may reasonably assume go unnoticed by the population. Note that unlike true polling error this bias is systematic, not random, and is designed explicitly to alter aggregate behavior while remaining within the noise. Advocacy in the model is dichotomous. Effectively, positive advocacy reports full global participation regardless of reality, and negative advocacy reports zero global participation regardless of reality.[7] As this should

[4] I focus on the interaction of the media with social network structure, rather than on coverage cycles as do Oliver and Myers (2003).

[5] This form makes clear how (relative) trust and attention come into play. When $M_S = .5$ both the media and social networks are equally influential. Increasing M_S simultaneously increases the weight on the media's information and decreases the weight on social network participation, so that relative trust in and attention to the media increase with increases in M_S.

[6] Assigning individuals' media strengths stochastically according to a uniform distribution yields similar results to setting all strengths at the distribution's mean of 0.5.

[7] Weaker messages are easily accommodated by varying the strength of the media, due to the fact that strength multiplies bias in the form of $c_{i,t}$. I explore this in analyzing multiple media sources.

on average reduce trust in the media, I focus on lower levels of media strength when exploring advocacy. Media outlets might turn to advocacy when, for example, it is sufficient for their needs, they possess charismatic advocates, or poll aggregators or other sources make apparent their earlier attempts at polling bias.

As noted in the text, bias in the model may substantively be as simple as alteration of reported polls, or as complex as the way in which issues are reported across the media. They all fit within the model as long as individuals proxy for them by considering the media to have provided a global participation rate biased relative to the true value in the population. Simple interpretations of bias might apply most usefully in the context of social movements and collective action, in which knowing how many others are participating can provide a sense of the worth of participating and, in some cases, the safety of doing so. More complex interpretations of bias, however, can capture different media intent. Bias capturing editorial slant or elite cues, passed along by the media, that try to push the population to support some candidate would seem to have similar poor intent to bias misrepresenting participation in some collective action or social movement. Bias due to sensationalism, in contrast, is a bit more neutral in intent: the media may just need more viewers. And bias can also arise when the public is misinformed about an issue and the media provide accurate information. This is still bias in the model because the media message differs from the true global rate of participation, but the intent is solely positive. We might think that this sort of "bias" occurs more or less often based on the difficulty of the issue (Carmines and Stimson 1980). In other words, "hard" issues might experience more lag in the public's understanding.[8]

Order of Operations

Each realization of the model begins with: 1) the creation of a network containing N individuals; and 2) the assignment of net internal motivations b_i to each person, as described above. In each period after this, the following sequence repeats.

1. Individuals simultaneously update external motivations according to the equations given above.

2. People simultaneously choose to participate or not in a given period according to the rule above.

3. Participation rates are calculated, both population-wide (p_t) and locally for each individual ($lpr_{i,t}$).

If no individual has changed her participation status for fifty consecutive periods, the realization of the model ends and final data for that realization are recorded. After 200 independent realizations for a given model parameterization are completed, means and standard deviations of equilibrium (steady-state) participation levels over these realizations are calculated. All simulation data were obtained via a JAVA program coded by the author. Increasing the number of realizations had no effect on outcomes.

[8] I thank an anonymous reviewer for pointing this out.

2 Network Construction and Parameterization

All network ties are symmetric and constant throughout each realization of the model. The latter is valid as long as the pace of network formation is slow compared to the rate of behavioral spread. The former is valid for forms of influence that involve reciprocity. The following is a brief description of the networks in the typology and their associated parameters. See Siegel (2009) for detail sufficient for replication.

- **Small World**: The base is a ring in which people are connected to *Connection Radius* other individuals to both sides of them. This sets the average connectivity. A *Connection Radius* of 5 thus indicates a connectivity of 10. Each connection has a chance equal to the parameter *Rewire Probability* of being severed and reconnected randomly. This sets the number of weak ties.

- **Village**: The base is an array of equally-sized groups called villages or cliques that are of size *Village Size*; every possible connection within these is made. If the population does not divide evenly in this way, then any left over individuals are placed into a final, smaller village. Each individual also has some probability, called *Far Probability*, of being connected to another individual outside of her village. These probabilities, which determine the number of weak ties, are checked twice, so the true probability of any individual's being connected to a particular person outside her own village is equal to twice *Far Probability.*

- **Opinion Leader**: Each individual is assigned a number of ties, k, according to the distribution $p(k) \propto k^{-\gamma}$, and connected randomly to this number of people. The parameter γ thus determines the characteristics of the network, with smaller values corresponding to greater leader influence, as there are more leaders with greater individual connectivity. These networks are also known as Scale Free networks.

- **Hierarchical**: The base is a hierarchy determined by parameter *Expansion Rate* in which one individual is placed at the top, and each individual in the network is connected to a number of individuals below her equal to *Expansion Rate*, continuing until no more individuals are left in the population. Thus, while each level of the hierarchy before the last one contains a number of individuals equal to a power of *Expansion Rate*, the last level may have fewer than this if the total population does not divide appropriately. Higher values of this parameter tend to lead to more influential elites at first, but for sufficiently wide hierarchies the influence of the top becomes watered down, particularly when motivations are correlated with network position. Each potential tie between individuals within the same level also has a probability equal to *Level Connection* of being made. The higher this probability, the greater the influence of followers.

3 Methodology: Sequential Parameter Sweeping

Deriving rigorous results from a computational model with more than a couple of parameters requires care. The method of theory-driven parameter sweeping I use largely matches that described and justified in Siegel (2009). It differs only in the following respects. First, the number of parameters is greater in this model, requiring additional parameter regions to be identified. These are described below. Second, I employ a supplementary generalized additive model (GAM) analysis (Beck and Jackman 1998). This entails sampling uniformly over the entire parameter space, and then fitting a GAM with these simulation data via a spline (or other) scatterplot smoother. This approach reduces problems, identified in Siegel (2009), of non-linearity and non-monotonicity in linear regression approaches to analysis, in that a GAM returns fitted values of each (flexible) function of the independent variables. However, this approach does not avoid the problem of sampling appropriately in regions of potential non-linearity or non-monotonicity. Because for this model the GAM analysis is unenlightening, likely due to an inability to sample on the specific regions of nonlinearity exhibited in the parameter sweeps displayed in the text, I do not include results of this analysis here.[9]

Sequential parameter sweeping necessitates building the model in stages. At the first stage, comparative statics for one or two parameters are directly computed and compared to extant theory. If similarly-behaving regions of this one- or two-dimensional parameter space can be identified, then a second stage of the model is added containing one or two new parameters to be varied. A set of comparative statics is computed for these new parameters for each of the regions identified in the first stage; the comparative statics on second-stage parameters hold for all first-stage parameters in a given region. This process can continue until no regions can be identified to simplify the analysis. For this model, the stages are: 1) aggregate participation without networks or a media; 2) participation in networks without a media; 3) participation in networks with a single, unbiased media outlet; 4) participation in networks varying bias (polling bias and advocacy); 5) participation in networks varying the number of biased media sources; and 6) participation in networks under selection into media.

The first stage corresponds to the parameter space spanned by $\{N, b_{mean}, b_{stdev}\}$. Siegel (2009) describes the analysis and theoretical support behind splitting this into three regions, and the text briefly discusses the relevant logic. As noted there, I call these regions the weak, intermediate, and strong motivation classes. They correspond to cases in which a population fully connected by a network (i.e., all possible connections are made) experiences a cascade of participation rarely, sometimes, and almost always, respectively. Higher values of b_{mean} increase participation levels in all regions. In line with limit theorems, increasing N reduces randomness in aggregate behavior, decreasing participation when it is unlikely, and increasing it when it is likely. This has the effect of reducing the likelihood of the intermediate motivation class. Increasing b_{stdev} increases participation as long as b_{mean} is not too high. The representative parameter triples used for plots of the intermediate and strong motivation classes are respectively $\{1000, .6, .25\}$ and $\{1000, .6, .3\}$. Note that these particular numbers are used for purposes of visual representation only; results discussed in the text as applying to motivation classes apply to all sets of parameters that fall into a

[9] The generalized additive models serve mainly as a robustness check. Additional simulation data beyond that included in this paper and appendix were taken and helped inform the analysis. Both additional data and the GAM analysis are available upon request. For further discussion of computational methodologies in general, see de Marchi (2005) and Miller and Page (2007).

given class.

The second stage is networks. I characterize each of the four network types according to the regions of the parameter space over which the model behaves similarly in response to the media. For purposes of visualization I choose representative parameter values for the figures. "Higher connectivity" lines correspond to a *Connection Radius* of 15 for a Small World network, or a *Village Size* of 25 for a Village/Clique network; "Lower connectivity" lines correspond to values of 5 and 5, respectively. The number of weak ties that is optimal depends on the connectivity parameter, so I list the values used in pairs. For the "optimal" Small World network, with the *Rewire Probability* second: $(5, .3)$. For the "superoptimal" Small World network: $(15, .7)$. For the "suboptimal" Small World network: $(5, .1)$. For "optimal" Village networks, with the *Far Probability* second: $(25, .004), (5, .003)$. For the "suboptimal" Village network: $(5, .001)$. "High Leader Influence" lines correspond to a γ of 1.4 in an Opinion Leader network, while "Low Leader Influence" lines correspond to a γ of 3.2 and "Intermediate Leader Influence" lines correspond to a γ of 2.2. Finally, lines for the Hierarchical networks depend on the parameter pair (*Expansion Rate*, *Level Connection*). "Narrower, Lower Influence Followers" lines use the pair $(10, .002)$, "Narrower, Intermediate Influence Followers" use the pair $(10, .007)$, "Narrower, Higher Influence Followers" use the pair $(10, .02)$, and "Wider, Lower Influence Followers" use the pair $(25, .001)$.

The third stage is media strength. I explicitly vary it in all the results in the paper. The fourth stage simultaneously varies M_S and $M_B = M_B(L) = -M_B(R)$ for polling bias. In all analyses in the fourth stage I did a full sweep over all parameters; however, for purposes of presentation only representative values of M_S were used. In Figure 4 of the text, these are $(.02, .12, .18, .32, .62, .78)$, and in Figure 5 of the text, these are $(.16, .40, .48, .60, .80, .98)$. Since the level of bias does not vary for advocacy, I show all values of M_S in Figure 6 of the text.

The fifth stage simultaneously varies M_S with both of the following parameters: M_L and $M_B = M_B(L) = -M_B(R)$ for polling bias. Absent selection opposing sources employing polling bias cancel out (see the text) and so there is no corresponding figure. Since the level of bias does not vary for advocacy, I show all values of M_L and use the following representative values of M_S in presenting Figure 7 of the text: $(0, .2, .6)$.

The sixth stage replicates the analysis of bias under selection. Figures 1-4 in Section 5 of this Appendix use the same parameters as Figures 4-7 of the text. Figures 5-10 display aggregate participation under selection minus the same absent selection, yielding positive values if selection increases participation. All parameters used in producing them are either listed on the plots or given earlier in this section. Figures 5-7 and 10 mirror 1-4, while 8 and 9 look at polling bias with two sources.

4 Measurement

This section provides some guidelines for measuring the model's parameters and primary dependent variable, the aggregate rate of participation. Though not necessary for understanding the model or its analysis, these can be helpful both in testing hypotheses drawn from the model and in conceptualizing some of the model's parameters. This section is meant to be read with the seventh section of the paper.

Participation in the model corresponds to the aggregate rate of support for a non-status quo option. It 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. Measurement of participation is best accomplished after surveys and other measures of the population have been taken and the action is fully realized. That is, measurement should focus on the final vote share or voter turnout, the total participation when the movement or protest has reached its peak participation, or the final support for each policy option at the time a decision over them must be made. This is so as to minimize issues of media bias in reporting aggregate behavior as well as the measurement of any transient states that may occur before the steady state is reached.

Table 1 provides a list of all model parameters and heuristics for measuring them. The first column specifies the broader parameter (e.g., motivation class, which represents the general proclivity of a population to participate). The second column specifies any relevant constituent parameters (e.g., internal motivation distribution mean and standard deviation and population size, all of which feed into motivation class). The third column details heuristics for the measurement of all constituent parameters. I elaborate on this table for the rest of this section. Note that I do not include information as to how to measure selection. As the model's implications are robust to the presence of selection, one can safely assume selection exists without needing to measure the media selection process. Note also that I do not describe how to identify network type: this was done in Table 2 in the seventh section of the text.

Parameter	Constituent Parameter	Measurement Heuristics
Media Strength	Rel Influence Media, Networks	Survey of Trust/Attention (Latent Variable)
	Rel Influence Media Outlets	
Media Bias	Advocacy	Direction of Constant Information
	Polling Bias	Poll vs. Average Poll/True Value
Motivation Class	Internal Motivation Mean	Survey of Support (Latent Variable)
	Internal Motivation St Dev	
	Population Size	Direct Measure
Non-elite (SW, V) Network	Average Connectivity	Number of Discussants
	Weak Ties	Discussant Commonality
Elite (OL, H) Network	Elite Influence	#, Stated Importance of Elites
	Follower Influence (H)	Subordinate Socialization Survey
	Elite Unity	Public Statements, Elite Survey

Table 1: Heuristics for Parameter Measurement

The model parameters most central to my analysis are those related to the media: *media strength* and *media bias*. There are two types of strength parameter in the model. The first is the primary media strength parameter, denoted M_S in Section 1 of this Appendix, which specifies the relative influence of the media compared to one's social network. This is sufficient for one media outlet or multiple unbiased media outlets. For two biased media outlets there is an additional strength parameter, denoted M_L in Section 1 of this Appendix, that specifies the relative influence of the pro-participation biased media outlet compared to the pro-status quo biased outlet.[10] All strength parameters can be measured similarly via survey instruments. Relative influence, the concept strength captures, is defined as the degree to which a particular source of information affects one's decisions (specifically, one's external motivations to participate) as compared to another source of information. It is an amalgam of trust in and attention to both sources of information, and the most important latent variable underlying survey responses to a battery of questions designed to measure trust in and attention to each of the information sources would be a good measure. This latent variable could be obtained via factor or principal components analysis, for example, or an item-response model (Treier and Jackman 2008).[11]

Now consider bias. There is one media bias parameter for every biased media outlet. Each describes the difference between the information put out by the media outlet and the aggregate participation rate of the population. Measurement of the bias of known advocates is comparatively straightforward. If the information from a media outlet does not vary over the course of the spread of participation (e.g., the information encourages support in a constant direction), then that outlet may be termed an advocate. Discerning the direction of an advocate's information (i.e., pro-participation or pro-status quo) on any particular issue is typically clear due to its lack of variation, but metrics based on word frequency might prove useful. All that is needed is the direction of the bias for an advocate in any case, so detailed metrics are not required to apply the model to one or more advocates.[12]

Polling bias presents more difficulty in measurement. When media reporting takes the form of simple polls, the degree of polling bias can be inferred by comparison with other sources of information. In declining order of accuracy, these sources might include the true population level of support, a polling average across many media outlets possessing biases in different directions, or a single baseline media outlet. Longer time series for comparison would improve the accuracy of the measurement by allowing one to weight different polls according to their historical accuracy, as do some poll aggregators. When media reporting is more complex, however, it can be difficult to quantify the specific effect of editorial slant, sensationalism, or even a desire to provide more accurate information than is available in the population. In line with this difficulty, there are numerous approaches designed to measure media bias but no single definitive one. The text discusses this point at more length. Sometimes natural experiments can be used, as in Druckman and Parkin (2005). But other times the best one might be able to do is a qualitative assessment of the direction of polling bias, which can still be useful information.

The next parameter is *motivation class*, which represents the intrinsic proclivity of a population to participate. There are three possible values of this parameter—weak, intermediate, and strong. Motivation class itself is determined by the joint action of three parameters: the mean and standard deviations of the distribution of internal motivations, and the number of individuals in the population. The size of the population is directly measurable. Very small populations are likely to fall into the intermediate class. Increasing the population can

[10] More biased outlets could be accommodated with more such strength parameters.

[11] Direct questioning about relative influence might also be effective, depending on the level of precision required.

[12] For example, it is generally not difficult to discern the direction of bias of Bill O'Reilly or Rachel Maddow from a cursory exposure to their shows.

move a population to either the strong or the weak class from the intermediate one. It also produces less aggregate participation in the weak class and more in the strong class. One needs to know the two distributional parameters to determine which class one is in. These are measurable as well, though less easily.

Each individual's internal motivation measures her proclivity to *participate*, where participate might refer to political participation or to supporting a new option over the status quo. This is best captured via a survey using either a single question that elicits nuance in proclivity or a battery of related questions. In the first case responses to the question directly make up the distribution of internal motivations. This is most useful when the survey question can be administered before the spread of participation or, failing that, early in the process of participatory spread (i.e., before external motivations have had time to play a large role). For example, the survey may ask about expectations to vote a year away from an election. In the second, one may derive a latent proclivity based on preferences over options related to but not the same as those over which the individual must now make a decision. In either case, the distribution of the population can be inferred from the sample distribution in the usual manner. Increasing the mean and the standard deviation tend to move the population toward a higher class (i.e., closer to the strong class) and increase aggregate participation within each class, since increasing each leads to more individuals with high internal motivations who play an outsize role in the spread of participation.[13]

A benefit of splitting the space of internal motivations into trichotomous motivation classes is that requirements of data quality are less than if one needed a continuous measure of motivation. Specifically, if one is only looking to discern whether the population is in the strong or intermediate motivation class, for instance, then one need not asses the distribution of internal motivations as fully as described in the previous paragraph. Small populations are much more likely to be in the intermediate class. For larger populations, early participation can provide an estimate of the number of early movers. Since early movers in the model act due to high internal motivations, the more of them there are, the more likely the population is in the strong class. Thus many early movers may be a sufficient proxy for being in the strong class.

Finally, there are the network parameters. There are four different *network types*, each with its own parameters: Small World, Village, Opinion Leader, and Hierarchy. See the Table 2 in the seventh section of the text for heuristics for identifying each. The first two lack elites and are described by parameters representing the average connectivity and number of weak ties between individuals.[14] The second two have elites whose innate proclivities to participate may be uncorrelated with their elite status, or positively or negatively correlated in the sense that elites have uniformly high or uniformly low motivations to participate.

[13]This is uniformly true for the mean of the distribution of internal motivations. Increasing the standard deviation also produces more individuals with low internal motivations, but these tend to be of lower importance in determining the spread of participation, as long as the mean is not too high. This is due to the path dependent spread of participation in the network. The logic is similar to that given for selection with a single media outlet in the sixth section of the text.

[14]See Section 2 in this Appendix for a discussion of differences in the parameterization of all networks and the exact method of network construction in the model. Because this level of detail is unnecessary for empirical testing at the level I discuss in the seventh section of the text, I keep to the more general description of the parameters here.

Both elite networks have a parameter that specifies elites' influence; the Hierarchy also has a parameter that specifies followers' influence.

The structure of the network can be mapped via techniques such as snowball samples, or specified entirely in an experimental setting. However, one of the virtues of this model is that relatively poor network data are sufficient to utilize the model's predictions. See the Table 4 in the seventh section of the text for more on this. One can assess average connectivity and weak ties via survey instruments that elicit the number of discussants on a particular topic and the degree to which an individual believes one's discussants also communicate on the topic. In elite networks one can assess elite influence via survey instruments designed to get at the number and perceived importance of elites in decision-making. Follower influence can be ascertained similarly, with inquiries into the frequency of relevant discussions among others at one's level in a Hierarchy. Finally, one can measure elite unity via both public statements by elites and elite surveys of their preferences.[15]

5 Figures for Media Bias under Selection

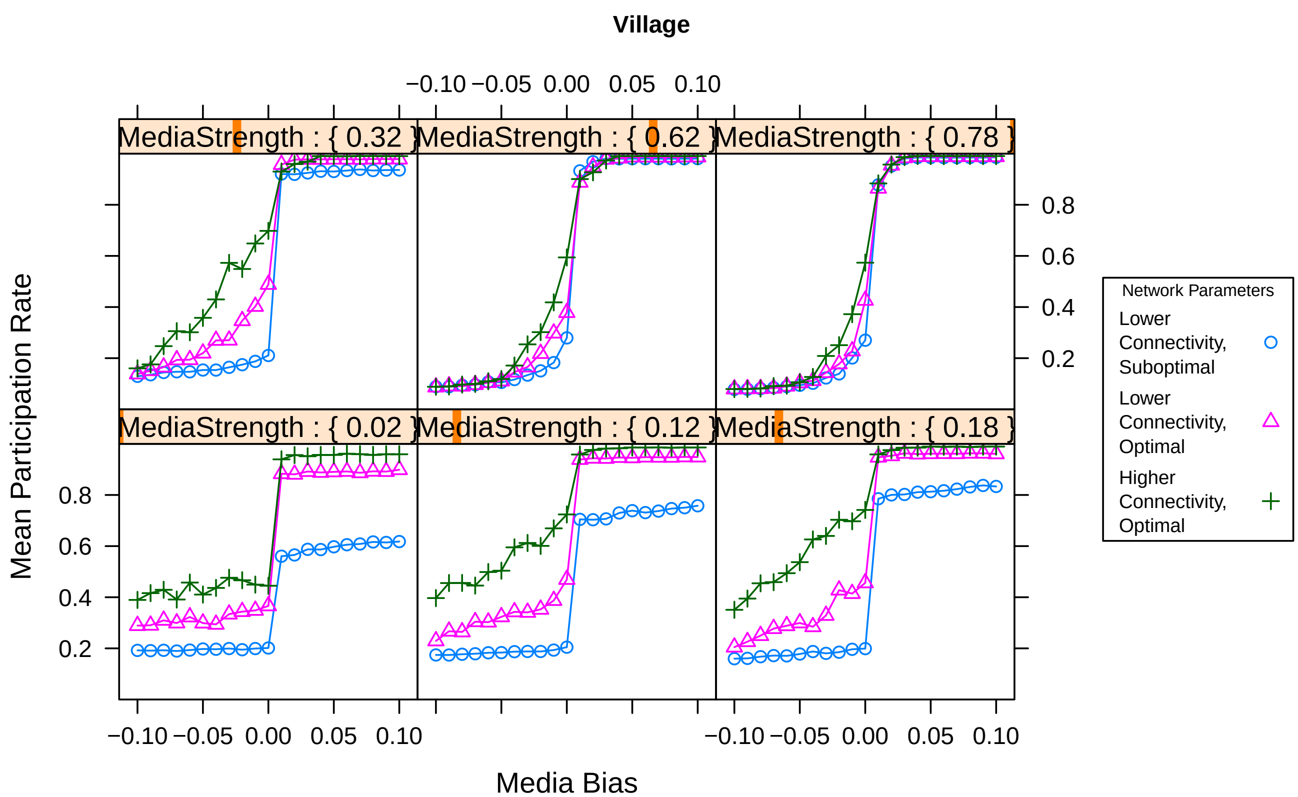

Figure 1: Polling Bias in Non-elite Networks under Selection

[15]The same technique as suggested to measure internal motivations could be applied to elites as well. In other words, fully measure elites while also sampling the overall population. However, since elites are comparatively little affected by others' behavior for much of the spread of participation, direct questioning of their preferences would likely produce less error than in the overall population.

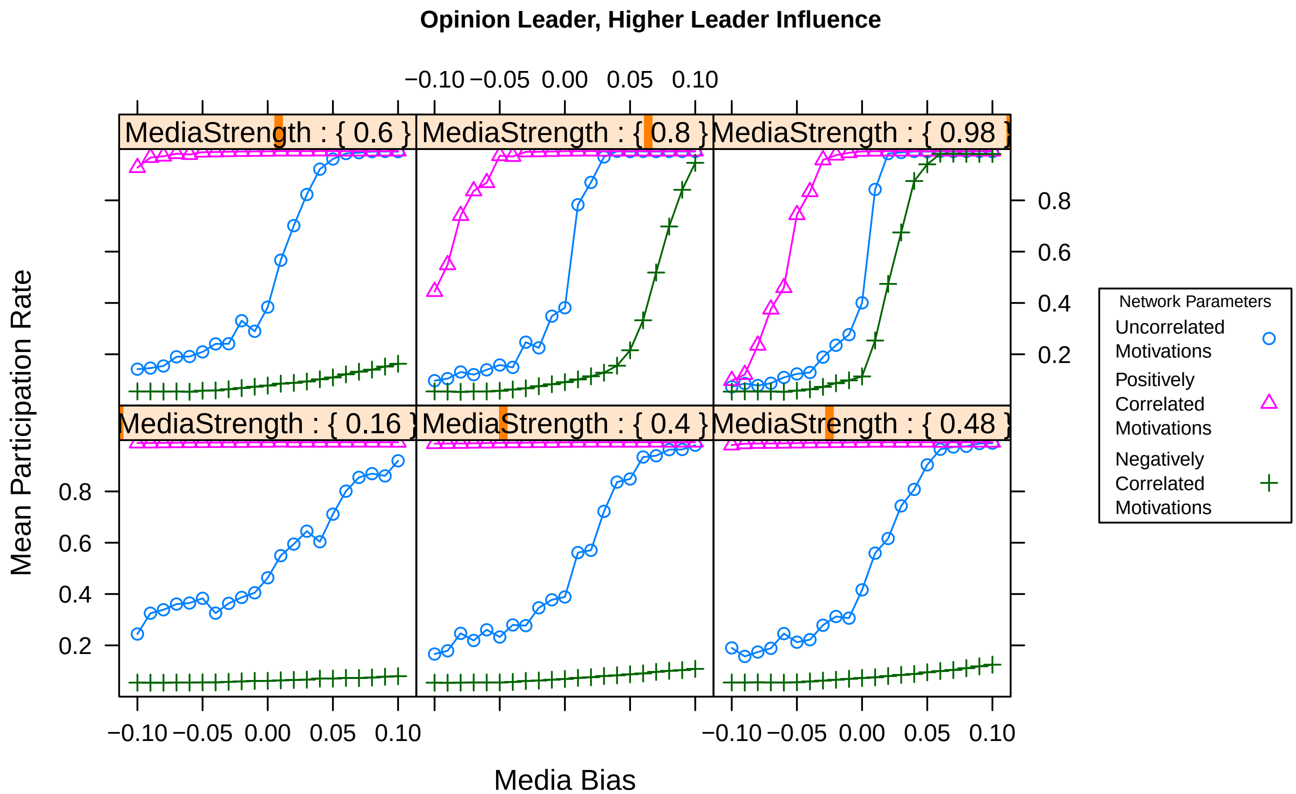

Figure 2: Polling Bias in Elite Networks under Selection

6 Media Penetration

The analysis thus far assumes that the media are a global information source that everyone receives. In other words, the media have full penetration in the population. In this section, I relax this assumption by varying the level of media penetration. All prior results hold under this relaxation, and one can derive additional claims as well. Chief among these is that increasing individuals' access to an informative media does not always spur participation. The reason for this is similar to that given for why increasing media influence does not always spur participation.

Media penetration in this section is a proxy for the level of access to the media. Full penetration implies that everyone has access to the media, while zero penetration implies that no one does. I use a simple index, M_P, to capture this idea. It ranges between zero and one, inclusive, and dictates the probability that each individual has access to the media. (I consider only baseline media here.) Access to the media for any individual is thus stochastically determined. I explore simultaneous variation of this parameter with media strength. Both act similarly in affecting outcomes in several respects.

I assume for this analysis that the media is an unbiased source of global information. Figures 11 and 12 show that media penetration acts in much the same fashion as does media strength. The series of six plots in each figure are conditional plots of the equilibrium participation rate against media penetration. The bar running along the top of each plot indicates the level of media strength for that plot. Figure 11 displays results for a Village network,

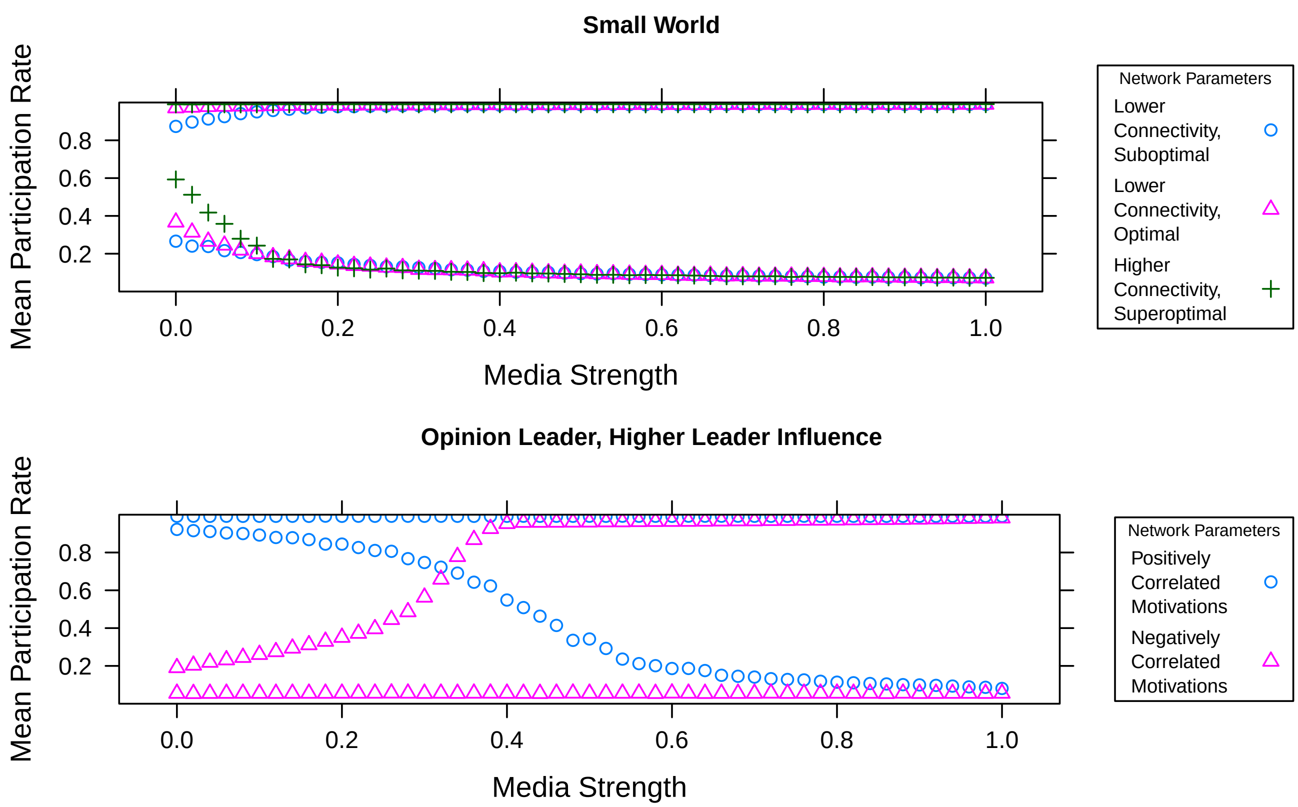

Figure 3: Advocacy under Selection

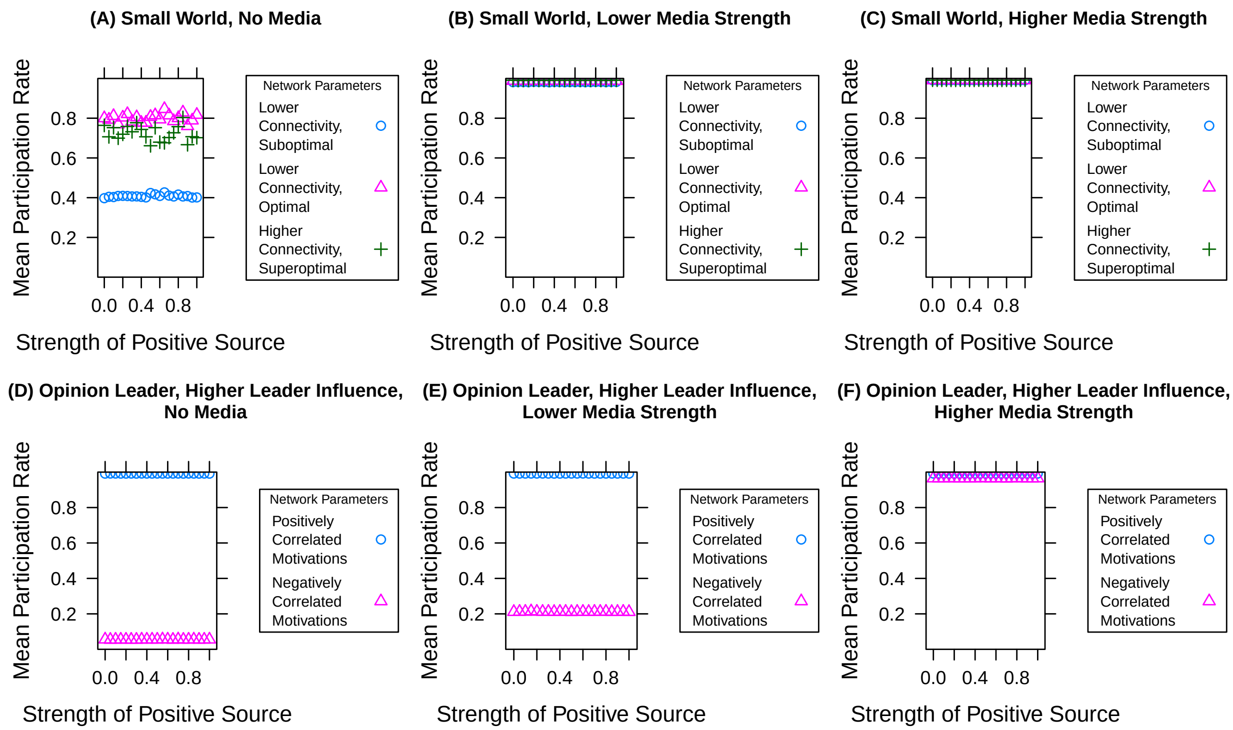

Figure 4: Advocacy—Two Opposing Media Sources under Selection

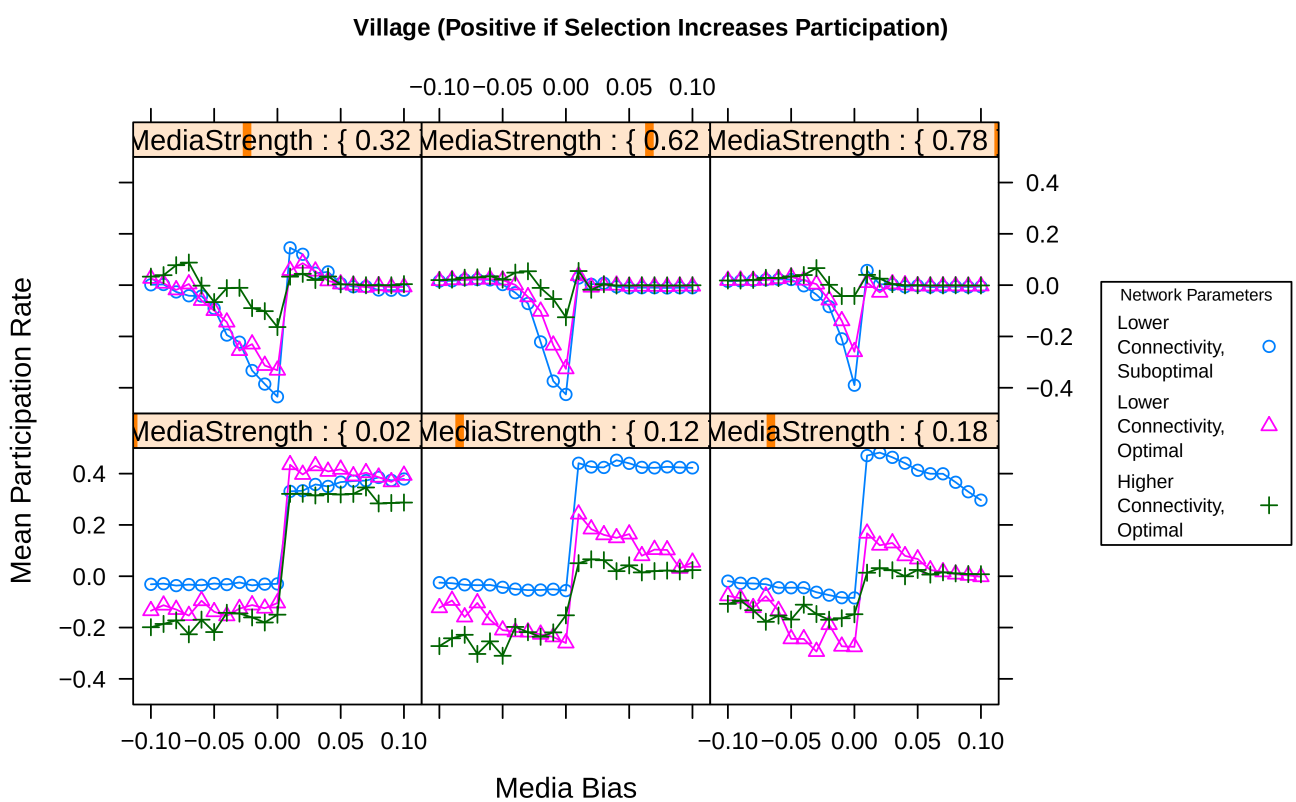

Figure 5: Comparative Effect of Selection in Polling Bias in Non-elite Networks

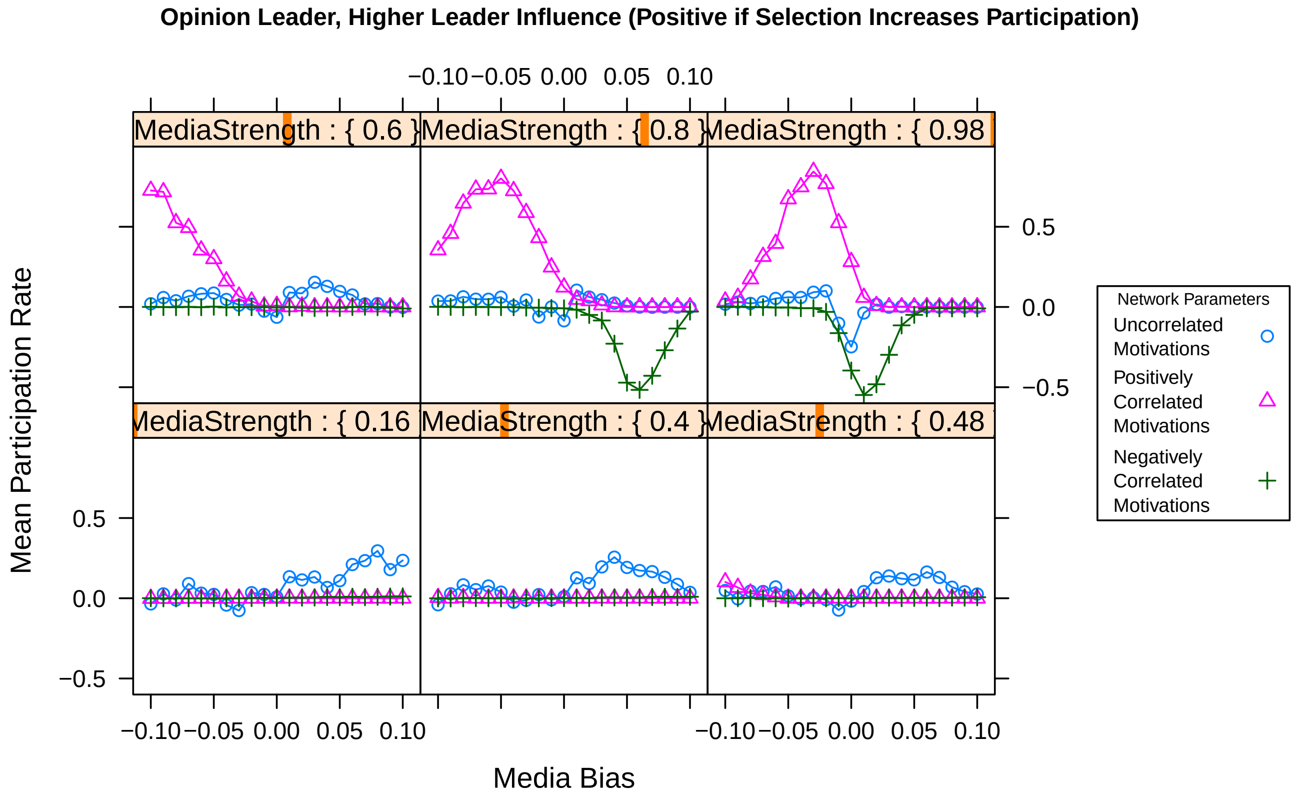

Figure 6: Comparative Effect of Selection in Polling Bias in Elite Networks

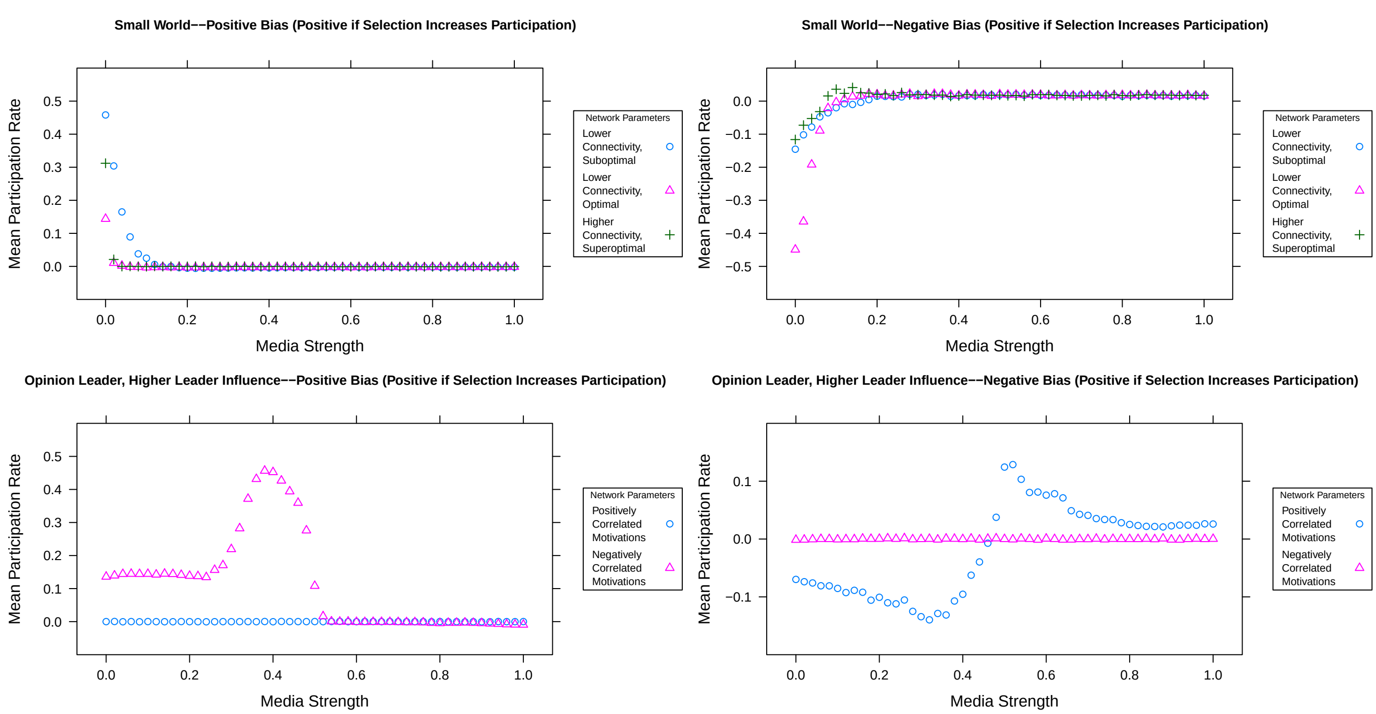

Figure 7: Comparative Effect of Selection in Advocacy

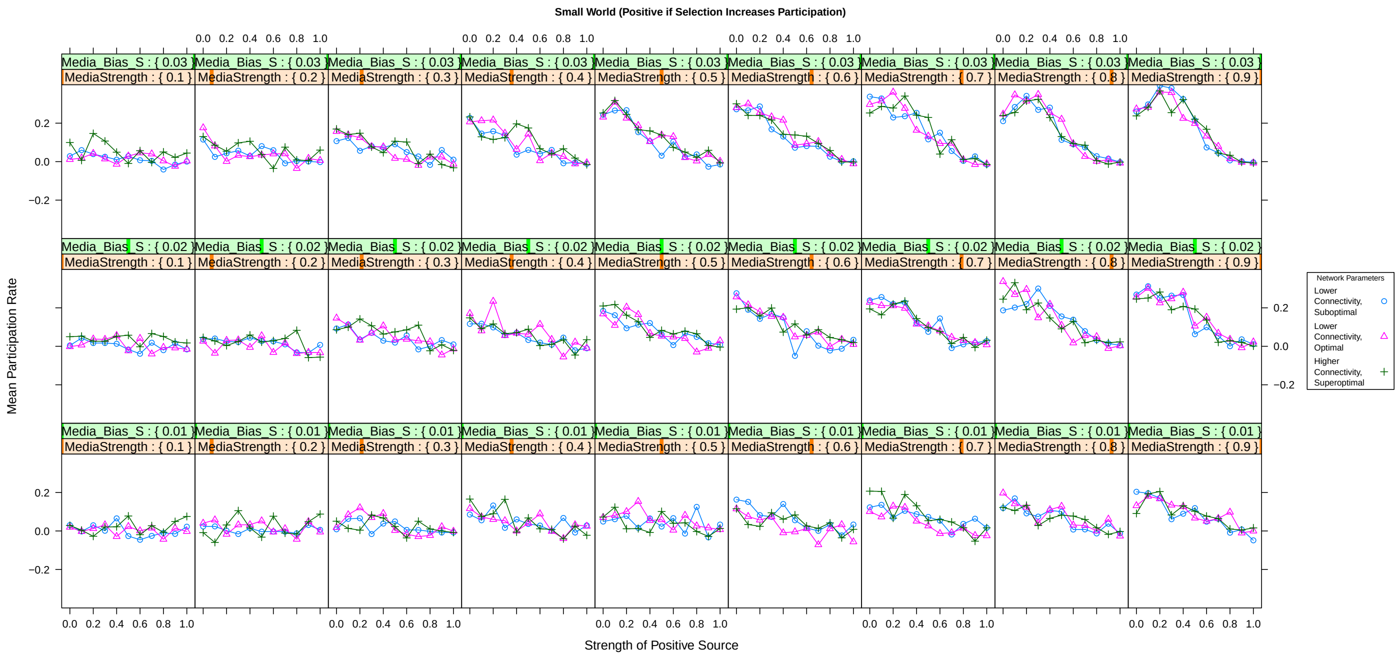

Figure 8: Comparative Effect of Selection in Polling Bias—Two Opposing Media Outlets, Non-elite Networks

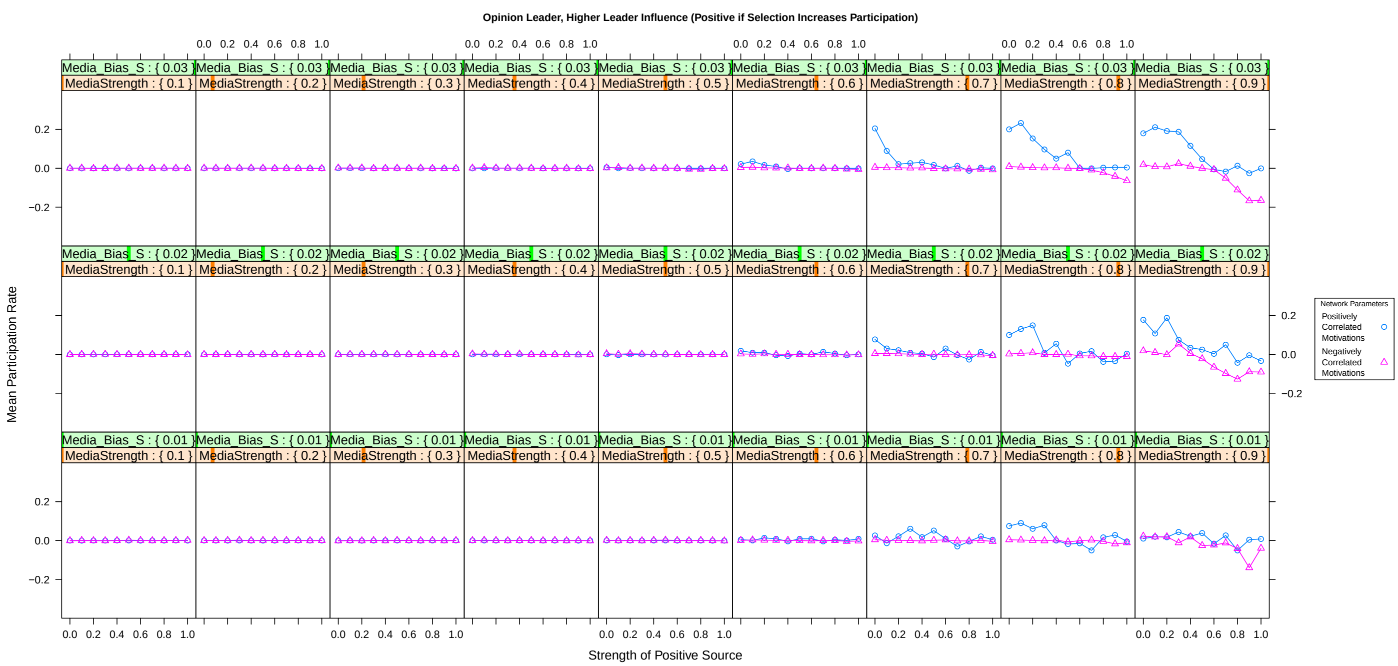

Figure 9: Comparative Effect of Selection in Polling Bias—Two Opposing Media Outlets, Unified Elite Networks

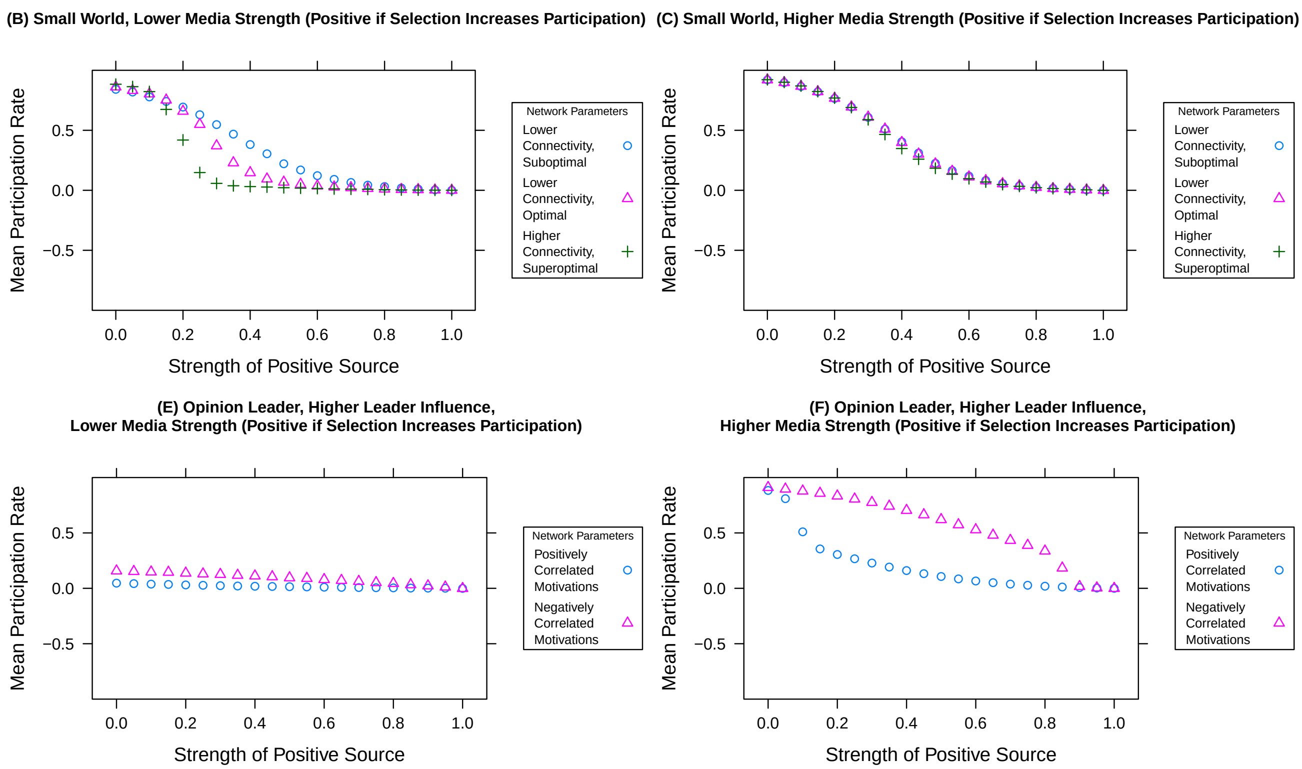

Figure 10: Comparative Effect of Selection in Advocacy—Two Opposing Media Outlets

Figure 12 for an Opinion Leader network. Different lines correspond to different network parameterizations, as usual. In Figure 11, the values of M_S are $(.02, .06, .12, .30, .46, .74)$, in Figure 12, $(.18, .40, .58, .68, .84, 1)$

The first thing to note in both figures is that increasing media penetration is likely to increase participation when media strength is low. At low levels of media strength the media plays a more encouraging role. Increasing media penetration at these levels of strength carries the spur of the media to more people, increasing aggregate participation. This generalizes: as long as increased media influence would increase participation, increasing the penetration of the media also increases participation.

Second, observe in Figure 12 that when elites are sufficiently numerous and unified, both the strength *and* the penetration of the media must be high for the media to play a role. Overcoming the pull of social elites requires not only influential media, but also media that reach nearly all of the population. External propaganda, even if entirely truthful, would seem to be less useful still in cases where there are strong, unified social elites, since it is likely the propaganda will not reach the entire population.[16]

Third, the top row in both figures indicates that, as with media strength, participation is non-monotonic in media penetration. That is, the optimum level of media penetration

[16] One might also imagine that the presence of strong, unified social elites might make high levels of trust in external propaganda unlikely. Ultimately this is an empirical question; if true then one would focus less on the plots corresponding to high levels of media strength in Figure 12 and so expect an even smaller role for the media.

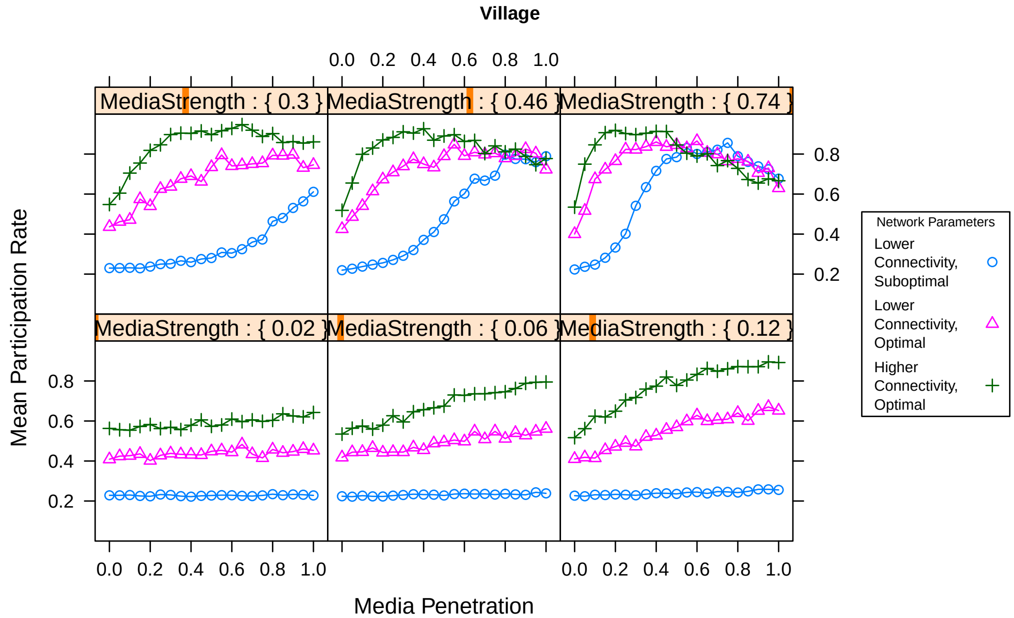

Figure 11: Optimal Media Penetration in Non-elite Networks

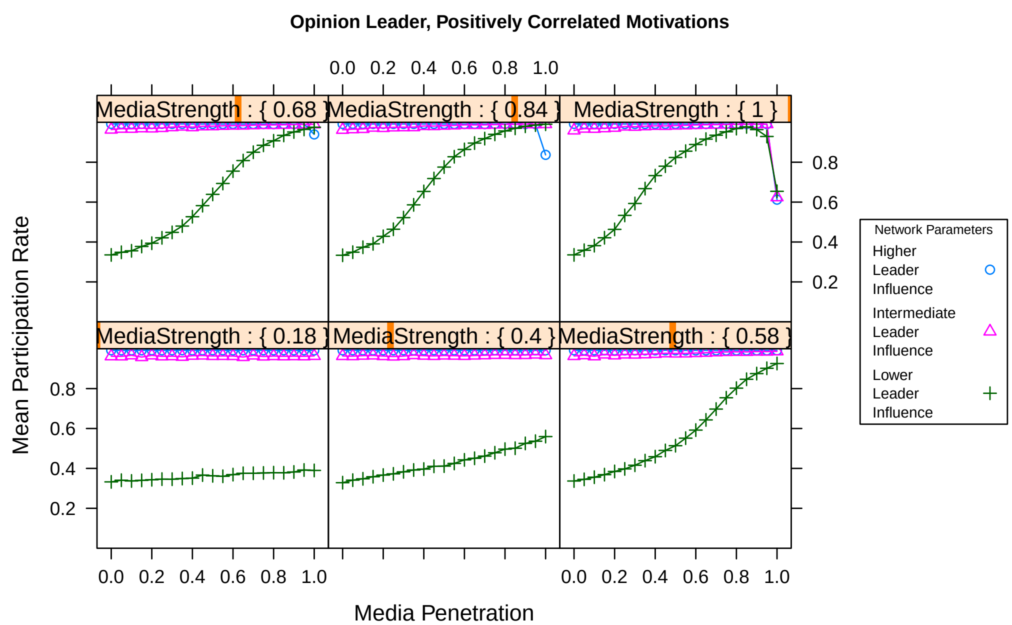

Figure 12: Optimal Media Penetration in Unified Elite Networks

occurs at less than full penetration. This too is a general finding: whenever increasing media influence would decrease participation, the optimal level of penetration is less than complete. In such cases, the diffusion of participation explicitly benefits from the unequal provision of the media. It is even possible to achieve marginally higher participation levels under unequal media provision than under full provision, for any level of media influence. This is due to the same trade-off discussed in the text: unequal provision of the media can provide less-connected, nonparticipating network clusters the information they need to spur their participation, while not destroying necessary enclaves of participation.[17] In the context of a social movement, then, it is not always the case that knowledge (of the global participation level) equates to power (in the form of a larger movement).

Finally, at a qualitative level, varying media penetration has a very similar effect as varying strength, as one can see by comparing the top right plot in Figure 11 to Figure 2C in the text. The former varies media penetration given a high-strength media, while the latter varies media strength given a media that reaches everyone. The dependence of participation on the independent variable is qualitatively similar in both cases, for each of the representative network parameterizations shown. This similarity in part made it possible to only analyze a media with full penetration in the text.

[17]Of course, given random media assignment this beneficial arrangement doesn't always happen. Sometimes the media is assigned to enclaves of participation, and not to the nonparticipating clusters, decreasing participation levels beyond what would have been observed if everyone had received the media. However, as shown in the section on multiple biased media outlets, the degree of encouragement occasioned by a beneficial arrangement of media access can outweigh the degree of discouragement by a detrimental arrangement, so that on average there is a positive, if marginal, effect of unequal provision on aggregate participation.

References

Axelrod, Robert M. 1997. *The complexity of cooperation: Agent-based models of competition and collaboration*. Princeton Univ Press.

Beck, N. and S. Jackman. 1998. "Beyond linearity by default: Generalized additive models." *American Journal of Political Science* pages 596–627.

Carmines, Edward G and James A Stimson. 1980. "The two faces of issue voting." *The American Political Science Review* pages 78–91.

Druckman, James N and Michael Parkin. 2005. "The impact of media bias: How editorial slant affects voters." *Journal of Politics* 67(4):1030–1049.

Granovetter, Mark S. 1978. "Threshold Models of Collective Behavior." *American Journal Of Sociology* 83(6):1420–1443.

de Marchi, Scott. 2005. *Computational and Mathematical Modeling in the Social Sciences*. Cambridge University Press.

Miller, John H. and Scott E. Page. 2007. *Complex adaptive systems: an introduction to computational models of social life*. Princeton, New Jersey: Princeton University Press.

Oliver, Pamela E. and Daniel J. Myers. 2003. "Social Movements and Networks: Relational Approaches to Collective Action." In *Networks, Diffusion, and Cycles of Collective Action*, Mario Diani and Doug McAdam, eds., pages 173–203, Oxford University Press.

Siegel, David A. 2009. "Social Networks and Collective Action." *American Journal Of Political Science* 53(1):122–138.

Treier, Shawn and Simon Jackman. 2008. "Democracy as a Latent Variable." *American Journal of Political Science* 52(1):201–217.