

Noise, bias, and expertise in political communication networks

Robert Huckfeldt^a, Matthew T. Pietryka^a, Jack Reilly^b

^a University of California, Davis, United States

^b New College of Florida, United States

ARTICLE INFO

Keywords:

Communication networks
Persuasion
Expertise
Opinion leaders

ABSTRACT

A central focus in the study of social networks and politics centers on the dynamics of diffusion and persuasion, as well as the manner in which these processes are affected by expert “opinion leaders.” The role of experts is particularly important in communication processes characterized by noisy, biased information – processes in which people with variable levels of expertise and strength of preference select informants, as well as being influenced by them. We employ an experimental approach that addresses these problems at multiple levels of observation in a highly dynamic context – small groups of individuals communicating with one another in real time. The analysis shows that participants formulate candidate judgments that decay in time, but the decay occurs at a significantly lower rate among the better informed. Moreover, the better informed are less affected by socially communicated messages regarding the candidates. Hence the influence of experts is not only due to their powers of persuasion, but also to the durability of their own privately formulated opinions. Their role in the communication process is further heightened by the higher value placed by participants on expert opinion, which in turn exposes the recipient to a heterogeneous and hence potentially influential stream of information.

© 2013 Elsevier B.V. All rights reserved.

Some individuals value political information as an end in itself (Fiorina, 1990), and hence the process of becoming informed generates intrinsic rewards, making the acquisition of information a self-reinforcing behavior. For others, the costs of information are so high that they swamp any benefit an individual might realistically expect to receive as a consequence of its acquisition (Wolfinger and Rosenstone, 1980; Downs, 1957). As a consequence, when left to their own devices, some individuals become politically expert while others remain politically naïve. In the spirit of Berelson et al. (1954) and Katz (1957), we should thus expect to see a division of labor in the communication of political information, with high cost individuals relying on others whose costs are minimal or even negative.

Complications arise because participants in the communication process are politically motivated, not only in their reasoning but also in their communication efforts (Kunda, 1999; Lodge and Taber, 2000). Hence socially communicated information is typically biased at its source, adding to the complexity of citizen decision making. Quite apart from these partisan biases, many of the underlying issues are imbedded in uncertainty, and even fully engaged individuals with shared political orientations might arrive at divergent political judgments. Thus individuals send and receive information that is not only noisy but also biased.

Within this context, Downs (1957) argues that an important way to minimize the costs of political participation is to obtain information on the cheap from other politically expert associates with shared political viewpoints. While this is a

seemingly efficient and reasonable strategy, its success is contingent on whether the supply of such informants is scarce or plentiful. In some contexts, the available informants might be either an expert with contrary preferences or someone with shared viewpoints who is bereft of useful information. In settings such as these, observational and experimental studies show that individuals often choose in favor of expertise (Ahn et al., 2010; Huckfeldt, 2001).

The resulting model of electorates – complex networks of interdependent actors with heterogeneous preferences and levels of expertise – raises a number of questions regarding the dynamics of becoming informed. How do individuals balance their own individually acquired information with information they receive from others? Do individuals evaluate new information in the context of old information (Lodge and Taber, 2000; Huckfeldt et al., 2004), or do they discard the old in favor of the new? Is the time-dependence of information and communication affected by individual expertise, by the reliance on socially mediated information, and/or by the heterogeneity of incoming information streams? What are the consequences of temporal dependence for the social diffusion of information?

We address these questions by constructing a small group experiment that implements variations in information costs across individuals, as well as making it possible for individuals to obtain information from one another. The experiment provides incentives for individuals to become informed, but these incentives must be assessed not only relative to information costs but also to the noise and bias attached to the information. Thus the subjects confront

challenges and dilemmas that parallel those faced by citizens in democratic politics.

1. Expertise, information costs, and interdependent citizens

Due to the individually variable costs of becoming informed, one might expect democratic politics to be driven by a cadre of self-appointed experts within the electorate – individuals for whom the problem of information costs are either greatly reduced, or for whom these costs do not apply. These experts are self-appointed because their roles are self-defined by their own interests and preferences in relationship to the value of political information. Having already paid the costs of becoming informed, the well informed are more likely to be politically engaged across a range of political activities, including the process of communicating their views to others (Huckfeldt and Mendez, 2008).

Such a view is premature for several reasons, and it runs the risk of exaggerating the net influence of single experts. First, the problem is not that experts are lacking in influence, but rather that many individuals receive multiple conflicting messages from experts with divergent viewpoints. Second, the recipients of messages are active participants in the communication process even when their supply of information is quite limited, and hence it becomes important to take into account the role of both the senders and recipients of information in the communication process. In particular, individual information processing strategies play a central role within communication networks, making it important to focus on the “nodes” as well as the “edges” – to address the role of individual recipients and communicators, as well the relationships that tie them to one another. This becomes particularly important relative to the value that recipients place on the information provided by alternative informants.

Downs' (1957) analysis assumes the importance of politically expert associates with compatible political orientations, but important problems relate to the identification and verification of an informant's expertise and trustworthiness (Boudreau, 2009; Lupia and McCubbins, 1998). Snowball surveys of naturally occurring communication networks show that individuals *do* communicate about politics more frequently with individuals whom they judge to be politically knowledgeable. Just as important, their perceptions of expertise among others are driven by the objectively verified expertise of potential informants. That is, they are typically quite accurate in recognizing the political preferences of those who are politically expert and engaged. These snowball surveys also show that the perceptions of expertise held by others, as well as the reported frequencies of political discussion, are only modestly affected by political agreement (Huckfeldt, 2001; Huckfeldt et al., 2000).

Moreover, when subjects in laboratory experiments are given the opportunity to obtain political information from other subjects, they place a greater emphasis on the expertise of other subjects rather than the presence of shared political preferences (Ahn et al., 2010). Similarly, in field experiments that address the natural formation of communication networks, both Lazer et al. (2010) and Levitan and Visser (2009) identify the minor role played by compatible political views in the formation of associational networks. In short, there is scant evidence to suggest that individuals effectively avoid any association with individuals holding preferences that are different from their own (see Huckfeldt and Sprague, 1995; Huckfeldt et al., 2004). Thus we turn to the role of the communication process itself to understand the manner in which noise and bias are filtered by the communication process within associational networks.

2. Memory constraints on the process of becoming informed

Time and the organization of human memory produce their own constraints on political communication and the process of becoming informed. Limitations on the capacity of working memory mean that individuals are continually storing and retrieving information in long-term memory, and information that is seldom retrieved becomes increasingly more difficult to recall. Time is certainly not the only factor affecting the accessibility of information from long term memory. Some information is more compelling (and hence retrievable) than others, both due to the inherent characteristics of the information, as well as to the correspondence between information characteristics and the cognitive map of the individual (Fazio, 1995; Berent and Krosnick, 1995). While time might play a potentially important and systematic role in the process, expectations diverge regarding the exact nature of the role, as well as the direction, of temporal effects.

First, as a counterfactual baseline, to the extent that individuals engage in memory based processing with infinitely accurate recall, the first piece of information obtained in reaching a judgment should be as important as the last piece of information. More realistically, to the extent that individuals engage in memory based processing with finite recall, we would expect a recency effect in which more recent information should have the greatest consequence.

Second, if the process of becoming informed is autoregressive (Huckfeldt et al., 2004), new information is processed in the context of old information. Hence, new information is less likely to be influential to the extent that it diverges from old information. In the context of memory decay, however, a persistent shift in the message being communicated ultimately swamps earlier signals in favor of more recent ones. In this way, an autoregressive process in the context of memory decay produces a complex moving average of messages, autoregressively upweighting earlier messages but simultaneously downweighting due to decay.

Finally, an on-line processing model employs an auto-regressive framework in which new information is judged in the context of old information (Lodge and Taber, 2000), but in this instance the effect of old information is summarized and consolidated in the form of a tally – an attitude or judgment that the individual brings to the interpretation of new information. When an individual receives new information in the on-line model, it is judged relative to prior judgments based on earlier information. In this case we see a primacy effect in which new information is less likely to be influential to the extent that (1) the pre-existent judgment is held more confidently and (2) the new information diverges from the old information. Here again, the primacy effect of earlier messages must compete with memory decay.

We rely on the early insights of McPhee et al.'s (1963) analysis in addressing the implications of social communication, political expertise, and memory decay for the political communication process. In his computer simulation, agents take information from sources in the environment, such as the media. They form prior judgments on the basis of that information and share their opinions with others. Based on these communications, they update these priors and communicate the information again. We pursue McPhee's contributions in the context of an experimental design and analysis that is inspired by a continuing stream of work in the study of social dilemmas.

3. The experimental design

Studies of political communication through social networks are beset by two related problems. First, social networks involve

explicitly endogenous processes. You choose your associates subject to contextually constrained supply, and then your associates influence you. Hence it is difficult to separate the influence of network construction from the influence of information transmission within and through networks.¹ Second, traditional experimental methods – the techniques which might be expected to resolve these problems – are difficult to implement in network studies that focus on communication as a process. That is, even the most ingenious laboratory studies that isolate individuals and study their responses to political stimuli will fail to approximate the repeated and influential social interactions underlying the complex communication processes and individual interdependence occurring within social networks.

Hence, rather than a single subject experiment, we turn to a different experimental tradition with earlier roots in small group research (Verba, 1961) and more immediate inspiration from the study of individual interdependence and social dilemmas (Fehr and Gächter, 2002). Similar to Fehr and Gächter's small group experiments with individual monetary incentives to study the implications of cooperation and defection for communication and the punishment of free-riders among interdependent actors, our experiment employs a small group experiment with individual incentives to address the implications of emergent communication networks for influence and expertise among interdependent individuals. Rather than an experimental paradigm in which a message or a signal is randomly manipulated and individual responses are examined, we experimentally manipulate political preferences and information costs to study the emergence of social networks and the communication process that these networks generate. In short, the study is designed to combine the advantages of small group dynamics with network representations of communication in the context of an experimental design.

The experimental setting is based on a mock election with two “candidates” who are not real human subjects, but are represented as positions on a one-dimensional policy space. The preference space varies from 1 to 7, where each participant has a unique integer position that remains constant across the periods in an experimental session, but candidate positions are reset at each period. The participant's goal in each period is to elect the “candidate” most closely matching her own position on the same dimension, and she is rewarded with a cash incentive if the closest candidate to her wins the election at that period. The exact positions of the candidates are not known to the voters, thereby creating an incentive to obtain information. Privately obtained information incurs costs, and these costs are also assigned randomly to participants. In order to minimize costs, participants have an opportunity to obtain free information from other participants, and to employ public information that is also free.

Seven subjects participate in each experimental session, where one subject in each session holds each of the positions from 1 through 7. Two subjects pay nothing for privately purchased information, two subjects pay 5 Experimental Currency Units (ECUs), and three subjects pay 20 ECUs. Verbal communication was not allowed during the experiment. Subjects were identified by their unique participant number ranging from 1 to 7. All decisions and information exchanges were made using networked desktop computers and z-Tree software (Fischbacher, 2007). Subjects were not

able to match others' participant numbers to the true identity of other individual participants in the experimental lab.

One criticism of this design is that it is not actually experimental because it does not focus on first order effects. The experimental treatment is implemented through the experimental manipulation of preferences and information costs, but the analysis is less concerned with first order consequences (i.e. how much information subjects purchase) than with higher order consequences. That is, we focus on the extent to which individuals with low costs come to occupy influential roles in the communication process, at the same time that those with higher costs depend more heavily on socially communicated expertise. Thus, according to this criticism, the analysis is actually based on a simulation rather than an experiment.

We appreciate this position, and we have no objection to readers who prefer to understand the design in this context. At the same time, we see no compelling or even desirable goal in narrowing the field of vision for experimental studies. Higher order consequences of experimental manipulations are still experimental effects, and excluding them from an experimental analysis seems both unwise and unnecessary (Ostrom et al., 1992; Ahn et al., 2009; McKelvey and Ordeshook, 1986).

4. The experimental procedure

Thus, each experimental session lasts for approximately 1 h, and includes an average of 9 periods. A new election with new candidate positions occurs at each period, but the subjects randomly assigned information costs and preferences are held constant for the entire session. Before an experimental session begins, participants are randomly assigned integer preferences and information costs that remain unchanged for the duration of the experiment.² Additionally, all participants are informed that Candidate A's position is between 1 and 6, while Candidate B's position is between 2 and 7. Then, in each of the approximately 9 periods per session, the following steps occur:

1. Participants receive 100 ECUs, of which 50 ECUs can be spent on information (hence subjects with an information cost of 20 ECUs can only purchase two “pieces” of information).
2. The two candidates' positions are drawn from the respective intervals.
3. Participants may purchase *private information* at their assigned cost.
4. After the subjects receive the information, they are asked to provide a prior judgment regarding each candidate's position, and they are truthfully told that their judgments will not be communicated to other participants.
5. A new computer screen shows each participant the preferences and the amounts of private information that each of the participants has purchased. Based on this information, subjects are allowed to make a first request for *social information* from one other subject. This request for social information is free to the sender and receiver. Potential informants are not required to comply with the request, and they are told that they need not provide the same information to all requestors. Participants almost always agree to provide information, consisting of a single message with information regarding each candidate position.

¹ Not all networks are endogenous to the choices of the participants, and indeed important experimental work has focused on the implications of exogenously imposed networks (Kearns et al., 2009; Ryan, 2011). Most work in political communication has addressed networks that are endogenous to individual choice, and that is the literature we address here. Our argument is that, even when individuals are given control over network construction, their choices are often constrained by larger social contexts and their own competing priorities.

² The relationship between information costs and preferences is established randomly as well, but it is held constant across experimental sessions. Hence every session has the following cost, preference pairings: 1, 20; 2, 5; 3, 20; 4, 0; 5, 5; 6, 20; 7, 0. For experiments in which the relationship between costs and preferences is varied experimentally see Ahn et al. (2010).

6. After receiving the information, subjects are asked to update their prior judgments – to offer a new judgment regarding the position of the candidate.
7. Steps 5 and 6 are repeated two more times. Hence subjects have the opportunity to make three information requests from other subjects, and they update their priors at each step. This produces a series of four judgments regarding the candidates' positions: a prior judgment after purchasing private information but before communication, as well as three updates after each of three communications with other participants. It is important to emphasize that the subjects are never provided with a summary of the information they have received – they assess and evaluate the information as it becomes available. All information is provided sequentially and incrementally, and the subject's challenge is to integrate and assess the information.
8. After communication is completed and subjects record their last updated prior, participants are provided a final opportunity to purchase a final piece of information at a cost of 10 ECUs.
9. The outcome of the election is revealed to the voters. If the winning candidate's position is closer to a voter than the losing candidate's position, the voter earns 50 extra ECUs. If the winning candidate's position is farther away from the voter's position than the losing candidate's position, 50 ECUs are subtracted from the voter's account. If candidates are equally distant from the voter, the voter neither gains nor loses. A voter could thus earn as much as 150 ECUs in a period, but only if she did not purchase any information (or if her information cost was zero). The minimum payoff is 0 ECUs – when a voter spends 50 ECUs on purchasing information and her candidate loses the election.
10. Participants are informed of their net earnings, which accumulate across periods.
11. Candidate positions are reset, and participants proceed to the next period. At the end of the experimental session, subjects are paid the show-up fee plus their total earnings in cash, where 100 ECUs equals one U.S. dollar. The range of total earnings, including the show-up fee, is from \$8 to \$17, and mean earning is \$12.

The analysis of this paper is focused on the process through which individuals formulate their final updated judgments regarding the positions of candidates. Hence we only consider the process through the seventh step described above, at which point the participants offer their assessments of the candidates. Our concern is with the culmination of the social communication process, thereby delaying an analysis of the implications for participants' vote to another effort.

The participants thus have three potential sources of information on which to base their judgments regarding the candidates. First, the public information that the two candidates' positions are drawn from different intervals could potentially help a voter in the absence of other forms of information.³ Second, voters are allowed to purchase unbiased but noisy private information on candidates' true positions. Third, each participant has an opportunity to request social information from other participants – information that is both noisy and potentially biased. That is, the requestor not only depends on the reliability of information that serves as the basis for *the informants'* judgments, but also on the ability and

willingness of the informant to compile and provide the information in an unbiased manner.

The proximate consequences of the experimental manipulations meet our expectations. Participants with higher costs obtain less private information, and participants who purchase more private information are better able to make informed choices. Mean information purchases are 2.8, 1.9, and 1.2 for subjects with costs of 0, 5, and 20 ECUs. Simple regressions of subjects' *final judgments* regarding candidate A's position on the candidate's *true* positions produce slope coefficients of 0.64 ($t = 13.6$) for those who purchased 3 or 4 pieces of information, 0.54 ($t = 9.4$) for those who purchased 2, and 0.25 ($t = 4.7$) for those who purchased 0 or 1.

Our interests in this paper reach beyond these first-order consequences, however. The communication process is complex, based on interdependent actors, and participants cannot assume perfect candor in the process. In the spirit of Downs (1957), Festinger (1957), Berelson et al. (1954), Katz and Lazarsfeld (1955), and others, we expect the process to be contingent on the preferences and expertise of informants, the range of available informants, and the potential for biased and misleading communication.

Finally, providing subjects with endowments and charging them ECUs for privately purchased information is, of course, an inexact parallel to real world information costs. In reality, information costs are typically realized in the form of opportunity costs. A sports fan might prefer to spend his time reading the sports pages, but some are willing to pay the costs of becoming informed by reading the opinion page as well, even at the cost of less time to read all the box scores. We are not pretending that the advantage of our experiment lies in its external validity. Indeed, our experimental participants are part of an experiment that only simulates a real world process, and the advantage comes in enhanced internal validity – an opportunity to make inferences regarding the consequences of variable information costs for network construction, political communication, and interdependence.⁴

5. Heterogeneity and bias within networks

In the context of Downs' analysis, the experimental participants should select *well informed* informants who *share their preferences*. The problem for individual subjects is that the supply of such informants may be limited. Each of the directed graphs in Fig. 1 illustrates one period (or election) within the experiment. The arrows point toward the individual from whom information is being requested. The size of the nodes reflects the amount of each individual's investment in private information, and it becomes clear that the more highly informed participants receive more requests for information.

At the same time, Fig. 1 also shows that participants must often choose between expert informants with preferences that diverge from their own and non-experts with preferences similar to their own. While individuals might prefer to have expert informants who share their preferences, their choices are limited by availability in their local contexts, with important implications for network heterogeneity and the communication of bias. We begin the analysis by examining the first order effects of our experimental manipulations on the creation of communication networks, the centrality of particular informants within the networks, and the potential for heterogeneity and biased communication that is produced.

³ The value of the public information should not be overstated. The interval boundaries on candidate positions overlap significantly, and hence there is no guarantee that Candidate A lies to the left of Candidate B. In this way the election is more like a primary election within a party rather than a general election contest between parties.

⁴ Do participants understand the process within which they are participating? While we do not debrief the participants after every session, we pre-tested the experiment to make sure that participants understand the procedure. Moreover, at the beginning of every session we include a practice period for instructional purposes. Finally, we carefully monitor the experiment, and it became clear that the participants understood the experimental process and procedures.

Table 1
Proximate effects of experimental manipulations.

A. Creation of dyadic ties: information request by amount of information purchased by the potential informant and absolute distance separating preferences within the potential dyad (logit; standard errors corrected for clustering).								
	1. All choices		2. 1st choice		3. 2nd choice		4. 3rd choice	
	Coef.	t-Value	Coef.	t-Value	Coef.	t-Value	Coef.	t-Value
Preference distance ^a	–0.14	3.56	–0.12	2.50	–0.17	3.49	–0.13	2.34
Information purchased	0.39	8.15	0.58	8.58	0.40	6.23	0.24	4.12
Previous information request	–2.20	7.08			–2.20	5.49	–2.37	7.79
Constant	–1.82	14.83	–2.51	14.67	–1.74	11.20	–1.26	6.76
N=	13,482 (84 subjects)		4494 (84 subjects)		4494 (84 subjects)		4494 (84 subjects)	
χ^2 , df, p=	85, 3, 0.00		74, 2, 0.00		50, 3, 0.00		64, 3, 0.00	
B. Predicted probabilities of requesting information by distance separating preferences within dyad and informant information.								
	Minimum information		Maximum information				Δ	
1. Any choice								
Minimal distance		0.12		0.28				0.28
Maximal distance		0.07		0.19				0.19
Δ		0.05		0.14				
2. First choice								
Minimal distance		0.07		0.42				0.35
Maximal distance		0.04		0.28				0.24
Δ		0.03		0.14				
3. Second choice								
Minimal distance		0.13		0.42				0.35
Maximal distance		0.06		0.28				0.24
Δ		0.07		0.14				
4. Third choice								
Minimal distance		0.20		0.39				0.19
Maximal distance		0.12		0.25				0.13
Δ		0.08		0.14				
C. Network centrality by information purchased and mean distance of preference from others in network. OLS regression w/standard errors corrected for clustering.								
	Coefficient		t-Value					
Mean preference distance from others in network from others in network		–0.25		1.49				N = 749 (84 subjects)
Private information purchased		0.73		10.17				R ² = 0.29
Constant		2.33		5.00				Root MSE = 1.39
D. Estimated bias ^b of message by the absolute difference between the preferences of the sender and the receiver. OLS regression w/standard errors corrected for clustering.								
	Candidate A		Candidate B					
	Coefficient	t-Value	Coefficient	t-Value				
Difference in preferences	0.15	3.33	0.15	2.93				
Constant	0.68	4.86	0.73	3.90				
N=	715 (76 subjects)				715 (76 subjects)			
R ² =	0.02				0.02			
Root MSE=	1.40				1.51			

A one standard deviation increase in preference divergence (0.58) produces a reduction in the predicted number of information requests by 0.145, and an increase across its range (1.5) yields a reduction of 0.375. In contrast, a one standard deviation increase in the amount of information purchased yields an increase of 0.88 requests, and an increase across its range (4) yields an increase of 2.92 requests.

^a Preference differences in Parts A and D of the table are measured as the absolute value of the difference within the dyad.

^b Bias is estimated as the absolute value of the distance between the message and the messenger's immediately preceding prior judgment regarding the candidates.

Criteria for selecting informants. This problem is addressed more systematically in Part A of Table 1, where participant information requests are regressed on the amount of information each of the other participants requested, as well as the distance between the preferences within the relevant dyad – the preferences of both the potential recipient of information as well as the potential provider. The response variable equals one if the subject requested information from the dyad's alter in a given period and zero otherwise. The table displays the results for all three social information requests, first pooling these requests, and then for each request individually. Thus, in models 2–4 of the table, each row in the data matrix is a dyad and model 1 pools these observations and hence each row is at the dyad-request level. This structure means that each individual participant appears multiple times within the data set.

Hence we apply a clustering correction on the standard errors of the coefficients (Rogers, 1993; Williams, 2000).

We do not restrict participants from making multiple information requests from the same participant during the same period. While it is a relatively rare event, Table 1 includes a control for whether the subject previously requested information from the potential discussant. Thus, in the second request for information, the indicator variable equals one if the subject's *first* request for information was to this potential discussant. In the third request, the indicator variable equals one if the subject's *first or second* request for information was to this potential discussant.

Each model demonstrates statistically discernible effects for the difference in preferences within the dyad, for the amount of information privately purchased by the potential recipient of an




Fig. 1. Directed graphs for typical periods in an experiment. Size of node indexes amount of information purchased. Direction of edge signifies the participant from whom information is being requested.

information request, and for previous information requests from a provider in the same period. Participants are more likely to request information from other individuals who (1) hold preferences similar to their own and (2) have made personal investments in privately acquired information. The control for repeated requests confirms that they are relatively rare, underlining the consequences of a constrained choice set on the supply of informants.

Based on the estimates in Part A of Table 1, Part B shows the corresponding changes in predicted probabilities of information requests across the explanatory variables for respondents, with the dummy variable for previous information requests from a particular individual held constant at 0 or no previous request. The first model generates an effect for the information level of the potential informants that is substantially larger than the effect for

preference, but it becomes clear that this first model is an average across the three requests for information that vary systematically across the exogenous factors. The initial request is highly responsive to the potential informants' information levels, but the importance of the information level is increasingly attenuated for the second and third requests. In contrast, the importance of shared preferences stays relatively constant across the three choices, but its effect never exceeds that of the informant's information level.

Network centrality in a context of limited choice. The problem is not that the criteria of choice are changing, but rather that the range of choices becomes increasingly limited. The context of the experimental group imposes limits on the ability to implement Downs' advice – participants are unable to locate sufficient numbers of experts with shared preferences. Network formation is thus subject to the constraints imposed by the particular configuration of the surrounding context (Huckfeldt, 1983; Huckfeldt and Sprague, 1995). Within this context, perhaps one of the most surprising results of the paper's analysis is that participants place a higher value on expert information, and hence they confront a heterogeneous stream of information.⁵

This result carries important implications for the structure of the communication network. In particular, it points toward experts as being particularly influential in the communication process, with high levels of network centrality. In Part C of Table 1, centrality is defined within a period (or election) as a subject's "indegree" – the number of requests for information received from other participants.⁶ Participants with higher levels of indegree are more central to the communication of information within the process – the information they communicate occupies more space within the communication process. When this measure of indegree is regressed on the absolute distance between preferences within dyads and the amount of information that a potential informant has independently acquired, it becomes clear that expertise trumps shared preferences as the most important factor explaining centrality.⁷

Implications for bias. The fact that participants are more likely to weight expertise over shared preferences in the selection of an informant produces obvious advantages. At the same time, it also has the consequence of exposing recipients to messages from politically divergent sources – messages that are more likely to contain biases introduced by the informant. The participants in our experiment send messages aimed at persuading the recipient, and hence the messages are contingent on the sender's goals. Participants are free to send different messages to different recipients, and messages typically carry a bias that is distinctive to the position of the recipient relative to the sender. Hence these messages are not unlike those frequently sent in ordinary political communication, where lively conversation is both informative and strategic.

Part D of Table 1 estimates bias in the messages sent by participants, where bias is defined as the distance between the message sent at the first opportunity for social communication during a period and the sender's immediately preceding prior judgment regarding the candidates. And this measure of bias is regressed

⁵ At the same time, it is important to recognize that the motivation to acquire information from an individual with shared preferences does not necessarily contradict the motivation to acquire information from an expert. One would have more confidence in the capacity of an expert with shared preferences to provide reliable information supporting those preferences.

⁶ We report an OLS regression, but the substantive conclusions are unchanged when employing a negative binomial regression.

⁷ A one standard deviation increase in preference divergence (0.58) produces a reduction in the predicted number of information requests by 0.145, and an increase across its range (1.5) yields a reduction of 0.375. In contrast, a one standard deviation increase in the amount of information purchased (1.21) yields an increase of 0.88 requests, and an increase across its range (4) yields an increase of 2.92 requests.

on the absolute distance separating the preferences of the sender and the receiver of the message.⁸ The regression shows that bias increases as the distance between the preferences in the dyad increase. The maximum distance between preferences is six units, and hence the maximum predicted effect on a single candidate message is 90 percent of one unit (6×0.15). While these are relatively subtle effects, they are not without consequence, and even minimal levels of bias can be consequential when candidates converge. In short, the participants must take into account the potential for bias as well as the inherent noise that accompanies information taken from a stochastic distribution, and hence it becomes important to address the process through which network effects are realized.

6. Information, memory decay, and autoregressive processes

Autoregressive models suggest that new information is judged relative to previously obtained information. Thus, new information that deviates from expectations based on past information would yield a diminished effect (Huckfeldt et al., 2004). We evaluate such a model in Table 2, which considers the subject's final summary judgment regarding a candidate's position as a function of (1) all three social messages, (2) the deviations of these social messages from the subject's judgments immediately prior to receiving the message, and (3) the interaction between the two. The models also include the individuals' original prior judgments based on individually purchased private information, the amount of private information purchased, and the interaction between the prior and the amount of private information purchased.

This model provides support for an autoregressive influence model. In general, the ultimate effect of a message is attenuated by the absolute size of its deviation from the subject's immediately preceding, contemporaneous judgment regarding the candidates. The first message fails to generate a discernible pattern of effects for candidates, but this is due to a pronounced pattern of decay in the messages' effects. More recent messages are consistently more influential because earlier messages tend to be forgotten (Lodge and Taber, 2000).

Fig. 2 graphically displays the autoregressive process for subjects' final judgments of candidate A. Part A of the figure demonstrates that the effect of the third message on the subject's final judgment of the candidate diminishes as the message diverges from the subject's previous judgment. When the message is within one unit of the subject's previous judgment, the model predicts that the subject's next judgment will move about 0.3 units in the direction of the message. In other words, if the subject previously believed the candidate was a 4 and the message suggests the candidate is a 5, the model predicts the subject's next judgment will be about 4.3. If the message is two units away from the previous judgment, its effect falls to about 0.2. Messages that deviate from the previous judgment by three units produce no discernible change on the final judgment. In other words, subjects dismiss such divergent information rather than integrating it into their judgments.

Part B of the figure displays the effect of a socially communicated message on the subject's final updated judgment when the message deviated from the subject's immediately prior update. That is, the figure illustrates the consequences of receiving a message that does not correspond to an individual's contemporaneous opinion. In particular, the figure focuses on the third social message and its deviation from the immediately preceding second update. When

the message suggests the candidate is on the far left (lower values), the effect of the deviation is positive. This positive effect counteracts the effect of the message which would otherwise pull the subjects' final judgment downward. Conversely, when the message suggests the candidate is on the right (higher values), the effect of a deviation is negative. Once again, this negative effect offsets the effect of the message, which would otherwise pull the subjects' judgments upward. Thus, the net influence of a message decreases with its deviation from subjects' immediately previous judgment.

In summary, this analysis shows three things. First, messages that are at variance with the subject's contemporaneous judgments at the time they receive the messages are less likely to be influential in affecting the subject's ultimate judgment. Thus, a subject's summary judgment is *not* based on a fresh look at all the evidence. Rather, the final judgment depends on the subject's preliminary judgments and the contemporaneous judgment at the moment when a message is received (Lodge and Taber, 2000). Second, and just as important, this analysis suggests a dynamic process in which more recent information is generally weighted more heavily than earlier information. Finally, it becomes clear that the experts form prior judgments that are much more likely to endure through this influence process, as well as to register stronger effects on their summary judgments. Hence we turn our attention to a simple model of the influence process to consider the consequences.

7. A simple model of the process

We begin by expressing the updating process for the subjects' judgments as a function of three factors: (1) decay in the most recently updated judgment, (2) decay in the initial (prior) judgment based on individually purchased private information, and (3) new incoming social information that is communicated by other subjects.

The effect of the prior. The model assumes that the initial (or prior) judgment, formed on the basis of privately purchased information, has an enduring effect that declines at a compound fixed rate between judgments. At the first update, the effect of the prior is wP_0 , where w is defined as $(1 - \text{rate of decay})$ and at the n th update, its effect is thus $w^n P_0$.

The effect of updated judgments. Updated judgments generate first order effects that also decline at a fixed rate. At the n th update, the effect of the previous update is αJ_{n-1} , where α is the survival of the previous judgment.

Incoming information. At the same time that the prior and the previously updated judgments are vulnerable to decay, the subject is responding to an ongoing stream of social information communicated by other subjects.

Hence the current judgment arises as a consequence of the rate of decay in an immediately prior judgment update, the rate of decay in an initial prior judgment, and the effect of contemporaneous social information.

$$\Delta J_t = -dJ_{t-1} + w^t P_0 + eI_t \quad (1)$$

where $\Delta J_t = J_t - J_{t-1}$; d = the rate of decay in the previous judgment, with d expected to lie between zero and 1; P_0 = the prior judgment based on privately purchased information; w^t = the effect of the prior at J_t , with w expected to lie between 0 and 1; I_t = incoming social information received at t ; and e = the educative impact of the new social information.

The model is rewritten as:

$$J_t = \alpha J_{t-1} + w^t P_0 + eI_t \quad (2)$$

where $\alpha = 1 - d$, the memory or survival of the previous judgment.

⁸ The number of observations in these models is lower than the previous estimates in Table 1 because they exclude the small number of subjects whose requests for information were denied.

Table 2
Final judgments by priors and messages, with messages contingent on contemporaneous judgments.

A. Candidate A	Coefficient	t-Value	
Initial (prior) judgment	0.16	2.38	
Prior × info. purchased	0.10	4.08	N = 749 (84 clusters)
Information purchased	−0.35	3.75	R ² = 0.58
First message deviation	0.06	0.71	Root MSE = 0.96
Second message deviation	0.20	2.46	
Third message deviation	0.50	6.38	
First message	0.13	2.75	
Second message	0.25	5.01	
Third message	0.43	8.81	
First message × deviation	−0.02	0.92	
Second message × deviation	−0.08	4.28	
Third message × deviation	−0.13	6.65	
Constant	0.18	0.63	
B. Candidate B	Coefficient	t-Value	
Initial (prior) judgment	0.13	2.08	
Prior × info. purchased	0.10	4.32	N = 749 (84 clusters)
Information purchased	−0.43	4.08	R ² = 0.59
First message deviation	0.12	1.53	Root MSE = 0.97
Second message deviation	0.44	5.42	
Third message deviation	0.54	4.93	
First message	0.16	3.21	
Second message	0.35	7.12	
Third message	0.44	7.20	
First message × deviation	−0.04	1.60	
Second message × deviation	−0.10	4.70	
Third message × deviation	−0.12	5.24	
Constant	0.28	0.81	

Initial (prior) judgment, the subject’s initial judgment regarding candidate positions, based solely on the public information and any private information the subject purchased.
Information purchased = the number of pieces of information the subject purchased at the beginning of the relevant period.

Message deviation = the absolute deviation between the message the subject received from another subject at the first request for information, and the subject’s most recent judgment regarding candidate position (for the first message deviation, the most recent judgment is the prior. For the second and third deviations, the most recent judgments are the updates following the first and second messages, respectively).

It is helpful to develop the model through recursion. At the first judgment (J_1) there is no previous judgment to update – only a prior based on private information plus new information, and thus

$$J_1 = wP_0 + eI_1 \tag{3}$$

Subsequent judgments update the immediately previous judgment as well as responding to the prior and new social information. Hence,

$$\begin{aligned} J_2 &= w^2P_0 + eI_2 + \alpha J_1 \\ &= w^2P_0 + eI_2 + \alpha wP_0 + \alpha eI_1 \end{aligned} \tag{4}$$

$$\begin{aligned} J_3 &= w^3P_0 + eI_3 + \alpha J_2 \\ &= w^3P_0 + eI_3 + \alpha w^2P_0 + \alpha eI_2 + \alpha^2 wP_0 + \alpha^2 eI_1 \end{aligned} \tag{5}$$

Pushing the model beyond the reach of our experimental observations yields

$$\begin{aligned} J_n &= (w^n P_0 + \alpha w^{n-1} P_0 + \alpha^2 w^{n-2} P_0 + \dots + \alpha^{n-1} w P_0) \\ &\quad + eI_n + \alpha eI_{n-1} + \dots + \alpha^{n-1} eI_1 \end{aligned} \tag{6}$$

To consider the long term dynamic logic, we take the equation to its limit. For n sufficiently large, the equilibrium is

$$J_n = \frac{w^n P_0 - \alpha^n w^0 P_0}{1 - \alpha/w} + eI_n + \alpha eI_{n-1} + \dots + \alpha^{n-1} eI_1 \tag{7}$$

Assuming that both α and w are bounded by 0 and 1, the effect of the prior converges on zero and the summary judgments inevitably depend on the continuing stream of incoming information, where the stream of information is weighted to favor the most recent information.

In short, the past is attenuated because this system of behavior forgets past events and past judgments rather than accumulating them – as any stable system must. How fast does the memory of this behavioral system decay? The key lies in the behavior of w^n and α^n . As α increases – as the immediately past updated judgment looms larger in the formulation of the current judgment – the importance of information received earlier maintains its effect longer.

Since the updated judgment is the mechanism whereby the prior is modified by new information, α also provides an index on the temporal durability of effects due to messages from other participants. As w increases, the importance of the prior takes longer to disappear. In this context, it is important to consider the dynamic implications in the short-term as well as the long-term, and hence to obtain estimates for the model parameters.

8. Estimating the model

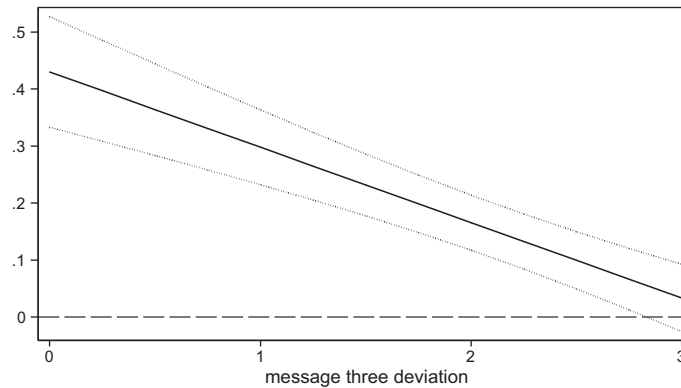
For purposes of estimation, we multiply both sides of Eq. (4) by α before subtracting the corresponding sides of the equations from Eq. (5). Upon rearrangement this yields,

$$J_3 = \alpha J_2 + w^3 P_0 + eI_3 \tag{8}$$


Hence regressing the final updated judgment regarding a candidate’s position on the previous judgment, the prior, and the incoming social information, provides statistical estimates for the model parameters – α , w , and e .

Part A of Table 3 displays the results of estimating the model in Eq. (8). For both candidate judgments, the final updated judgment (J_3) is regressed on the immediately preceding updated judgment (J_2), the initial prior judgment (P_0), and the immediately preceding (third) piece of communicated information (I_3). In view of the demonstrated importance of private information consumption, the

- A. Marginal effect (with 95% confidence bounds) of the third message on the subject's final judgment by the message's deviation from the subject's immediately prior update.



- B. Marginal effect (with 95% confidence bounds) of the third message's deviation from the immediately prior update by the third message.



Source: Estimates in Table 2a.

Fig. 2. The autoregressive influence of social information.

regressions also include the amount of information purchased as an explanatory variable, as well as its interaction with both the third installment of communicated information and the prior judgment.

First, the table shows that the initial (prior) judgment has no effect absent the interaction with the amount of information purchased. That is, the prior only matters among those participants who invest in private information, and the effect is enhanced by the level of that investment. Second, as would be expected, the table shows a substantial effect due to the immediately preceding update. Third, the model shows a substantial effect due to the final (third) message that appears to be at least modestly attenuated by the amount of private information purchased by the participant.

Part B of Table 3 adds an interaction between the third message and its deviation from the subject's immediately preceding (third) updated judgment. This yields no change in the estimates from Part A of the table, and it fails to produce a discernible effect due to message deviation. Hence we pursue the analysis based on the results in Part A. In view of the results shown in Table 2, the lack of an interaction effect due to the deviation between the message and its distance from the subject's immediately previous update warrants explanation. The difference is that the models in Parts A and B of Table 3 include updated judgments as regressors, while the models in Table 2 only include the prior, the stream of incoming messages, and the interaction with the divergence of these messages from participants' immediately preceding updates. These results thereby support the on-line processing model of Lodge et al. (1995) – the

memory of past information is mediated by past judgments. The information is not recalled directly and has no lasting effect, except as it forces updates in judgment.

Part C of Table 3 shows the estimated model parameters adjusted for the amount of information purchased by the subject. The results show that the survival of the prior is directly related to the amount of information purchased. Indeed, absent the interaction of the prior with the amount of information purchased, the prior has no effect. (This result means that we cannot reject the null hypothesis that w is 0 for the subjects who purchased no information.) In contrast, however, there is a dramatic effect of information investments on the survival of the prior. If we fail to reject the null that there is no effect among those who did not purchase information, w varies from 0 to 0.64 for the judgment regarding candidate A.

The effect of socially communicated information is also dependent on information purchases. Those who did not purchase any private information show an effect that is more than 3 times larger than the effect among those who purchased four pieces of information on each of the candidates. In short, those who purchase private information do not pay much attention to socially communicated information, and those who do not purchase private information are reliant on social communicated information obtained from other subjects.

What are the dynamic implications? Part A of Fig. 3 shows the decay in the prior over time for judgments regarding Candidate A. The figure plots the influence of the prior at the time the prior

Table 3
Estimating the dynamic model of judgment formation.

A. Final updated judgment by initial (prior) judgment, previous update, and final communicated information, with interactions.				
	Candidate A		Candidate B	
	Coefficient	t-Value	Coefficient	t-Value
Initial (prior) judgment	0.004	0.09	0.04	0.97
Immediately previous (third) update	0.66	11.81	0.63	10.17
Third message	0.21	4.88	0.21	3.55
Prior × info. purchased	0.06	2.82	0.07	3.29
Third message × info. purchased	−0.03	−2.28	−0.04	1.64
Information purchased	−0.08	0.93	−0.16	1.57
Constant	0.36	1.62	0.53	1.80
N=	749 (84 clusters)		749 (84 clusters)	
R ² =	0.69		0.70	
Root MSE=	0.81		0.83	

B. Final updated judgment by initial (prior) judgment, previous update, and final communicated information, with interactions.				
	Candidate A		Candidate B	
	Coefficient	t-Value	Coefficient	t-Value
Initial (prior) judgment	0.007	0.14	0.04	0.95
Immediately previous (third) update	0.65	10.23	0.63	9.60
Third message	0.25	4.19	0.21	3.52
Prior × info. purchased	0.06	2.74	0.07	3.27
Third message × info. purchased	−0.03	2.22	−0.03	1.54
Information purchased	−0.08	0.87	−0.16	1.63
Third message deviation	0.04	0.48	0.03	0.33
Third message deviation × message	−0.02	0.97	−0.002	0.09
Constant	0.36	1.41	0.47	1.95
N=	749 (84 clusters)		749 (84 clusters)	
R ² =	0.70		0.70	
Root MSE=	0.81		0.83	

C. Model parameters adjusted for individual information purchases, based on the estimates from Part A.										
Candidate A	Candidate A					Candidate B				
	Amount of information purchased					Amount of information purchased				
	0	1	2	3	4	0	1	2	3	4
W	0.16	0.40	0.51	0.58	0.64	0.35	0.49	0.57	0.64	0.69
α	0.66	0.66	0.66	0.66	0.66	0.63	0.63	0.63	0.63	0.63
e	0.21	0.17	0.14	0.10	0.06	0.21	0.17	0.14	0.10	0.06

is given (J_0) through the final judgment (J_3) for individuals purchasing four, two, and zero pieces of information. In general, we see quite rapid decay in the effect of the prior, even among those individuals who purchased the maximum of amount of private information.

In contrast, Part B of the figure shows the decay in the effect of socially communicated messages from the first judgment subsequent to their reception (J_1) through the final judgment (J_3). In this instance the decay occurs relatively more slowly, but in contrast to the effects of the prior, we see the greater effect of communicated messages on those subjects who purchased less private information. That is, among those who purchased four pieces of private information, their fourth and final judgment relies on approximately 40 percent of the prior, but is virtually independent of the first social message. In contrast, among those who purchased no information, the contribution of the prior has disappeared, and we see a modest effect due to the first social message.

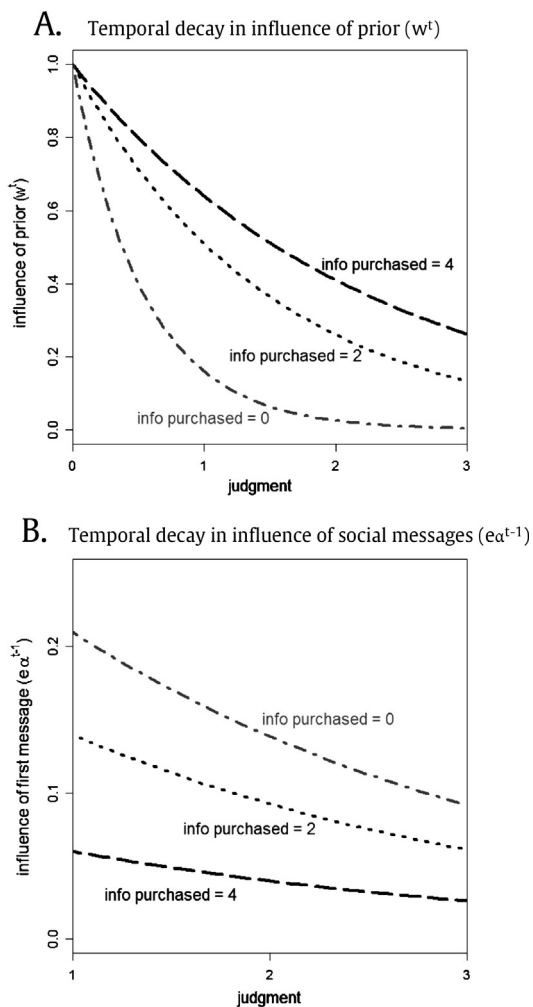
The implications are quite important. While experts play a central role in the process of political communication, we should not view them as wholly independent actors who emit signals but do not receive them. While experts are able to assess incoming signals in the context of their own accumulated knowledge, their judgments based on this knowledge decays with time, and this process of decay is offset by new, socially communicated information. As

Katz (1957) informed us more than 50 years ago, politically expert citizens are not immune to the effects of social communication.

Is the process autoregressive? The results show an interplay between recency and primacy in the communication process. Everything else being equal, recent communications matter more than earlier communications, and the decay of earlier communications enhances the relative effect of the most recent communications.

At the same time, decay is mediated by expertise. Individuals who invest more heavily in the acquisition of private information demonstrate a more enduring effect due to their priors, at the same time that they rely less heavily on messages obtained from others. As a consequence, they engage in on-line processing (Lodge and Taber, 2000), where new information is judged and assessed in the context of pre-existing judgments, and thus they tend to be more intransigent in their opinions.

In contrast, those who do not invest in private information rely less heavily on their prior judgments, and they pay more attention to new messages and new information. Hence the updating process takes on relatively more importance. The process is still autoregressive – informationally impoverished individuals do not exercise a comprehensive, memory based processing strategy. It is simply the case that, lacking a strong prior on which to formulate a judgment, they necessarily rely more heavily on contemporaneous information.



Source: Estimates in Table 3a.

Fig. 3. Implications of the model.

9. Conclusion

Political information is both noisy and biased as a direct and inevitable consequence of its subject matter. For many people in many circumstances, political issues are complex and ambiguous, with frequently high levels of uncertainty. At the same time, many other individuals are deeply invested in particular opinions and attitudes and hence care very deeply about political issues and outcomes. Communication regarding politics thus reflects both the complexity of the subject matter and the existence of strongly held opinions. Moreover, the attentive are being continually bombarded by heterogeneous streams of information that are both biased and difficult to validate.

In this context, citizen experts play an important role, even if most experts are self-appointed. In the context of citizen politics, the experts are those individuals who care enough about politics to pay the costs of becoming informed. For many, the costs are not so terribly high for the simple reason they thrive on the acquisition and analysis of political information. Whether the favored information source is the New York Times, the Daily Show, Fox News, or Rush Limbaugh, the expert is often ready and willing to pay the costs of becoming informed. Or alternatively, the experts encounter no costs because they enjoy the process of becoming informed.

The citizen experts are important for several reasons. They fill the social airwaves with political content. They package the information they transmit within inevitable patterns of bias that reflect

their own interests and concerns. They provide an information source for those political non-junkies who would just as soon spend their time in the pursuit of other interests and avocations. Most importantly, they inject political content into the everyday patterns of communication and interdependence that exist among citizens in democratic politics.

In some ways, political experts display high levels of political self-reliance. Through their private investments in political information, they develop their own political priors regarding a vast array of issues and problems. At the same time, this analysis also suggests that the independence and self-reliance of experts can be overstated. Knowledge and information are indeed fleeting, if only due to the inevitable processes of decay that undermine the certainty of almost any judgment. Hence, political junkies sustain their positions as experts, not by the breadth of their knowledge regarding things political, but rather by their continuing pursuit of political information. Their energy in the pursuit of politics gives rise to the continual formulation of new prior judgments regarding a great variety of issues and problems.

Moreover, this paper's analysis suggests that even experts are affected by patterns of communication with others, and sometimes these others may be less politically expert. Indeed, a defining ingredient of the expert is a willingness to engage in political communication with others, and our results show that this communication is not without consequence. Hence the influence of political communication may sneak in through the back door, integrating the expert within patterns of communication that rebound as a source of influence on the prime mover.

Finally, the analysis suggests that the process of becoming informed is autoregressive – individuals encounter and digest new information in the context of old information. This autoregressive process is best understood as an on-line process among the politically expert because these are the individuals with well developed attitudes, opinions, and beliefs. As we have seen, strongly held opinions survive the process of on-line updating, and hence the experts tend to be only modestly affected by new information. In contrast, the judgments formed by the less expert are not anchored in strongly held priors, and hence they are much more susceptible to messages received from others.

This research effort, like every other, involves advantages, disadvantages, and tradeoffs. Our research design serves as an abstraction from the real world at various points. We previously discussed the limitations on paid incentives as representations of information costs, but we employ other abstractions as well. Perhaps most importantly, while we provide an opportunity for participants to update their beliefs about the candidates, we do not provide an opportunity to update their opinions and preferences. That is, we assign fixed preferences for the discussants that lie beyond the reach of the communication process, but in reality, people's opinions and preferences are also subject to communication effects. At the same time, underlying political preferences are likely to be more stable than beliefs regarding candidates or the state of the political world.

In addition, this paper's analysis only deals with the search for political information, but this is only a minor part of daily life for most reasonable individuals. Indeed, the number of pure information seeking interactions is likely to be quite rare in most populations, although even nonpolitical interactions can provide incidental benefits with respect to political communication. In these instances as well as others, our effort is to gain leverage on the very real aspects of underlying social process that are difficult to observe, but these efforts must certainly be seen within the context of these simplifying abstractions.

A question that naturally arises is whether the social communication process we have specified constitutes a Bayesian updating process. That is, are the subjects in our experiment employing

Bayesian reasoning when they formulate prior judgments based on private investments in information and then update those judgments based on information that is socially communicated? If the process is Bayesian, the updated judgment should represent both the prior and the newly acquired information, where both are weighted inversely by their respective variances (Bullock, 2009; Bartels, 2002; Gerber and Green, 1999). Such an account runs at least partially parallel to the social communication process analyzed in this paper. Individuals form final judgments based on priors weighted by the amount of unbiased information they use in formulating the prior – an entirely reasonable (inverse) indicator regarding variance around the prior. Moreover, the analysis suggests that participants are cautious regarding the value of new information. They seek to minimize misleading bias by locating informants whose interests coincide with their own, and they are skeptical regarding new information, particularly information taken from non-experts whose judgments they do not trust.

The comparison to a Bayesian process thus provides an interesting frame of reference for evaluating the implications of social influence in politics. There is certainly no evidence to suggest that the participants in our experiment are self-consciously invoking Bayes' theorem. At the same time, they may be invoking standards of judgment which unintentionally approximate a Bayesian process. Alternatively, this social communication process might also be understood in terms of motivated reasoning. After investing more heavily in the formulation of their own prior judgments, the experts among our participants are personally committed to these judgments and less likely to be swayed by information to the contrary (Kunda, 1999; Lodge and Taber, 2000). Thus, in at least some instances, a process of motivated reasoning might produce a social influence process that parallels Bayesian updating. Under what conditions is a social influence process likely to take on the characteristics of Bayesian updating? While a full treatment of this question goes beyond the bounds of the current paper, the evidence points toward this being an important issue.

References

- Ahn, T.K., Mark, I., Tim, S., 2009. Coming and going: experiments on endogenous group sizes for excludable public goods. *Journal of Public Economics* 93, 336–351.
- Ahn, T.K., Huckfeldt, R., Ryan, J.B., 2010. Communication, influence, and informational asymmetries among voters. *Political Psychology* 31, 763–787.
- Bartels, L.M., 2002. Beyond the running tally: partisan bias in political perceptions. *Political Behavior* 24, 117–150.
- Berelson, B.R., Lazarsfeld, P.F., McPhee, W.N., 1954. *Voting: a Study of Opinion Formation in a Presidential Campaign*. University of Chicago Press, Chicago.
- Berent, M.K., Krosnick, J.A., 1995. The relation between political attitude importance and knowledge structure. In: Milton Lodge, Kathleen, M., McGraw (Eds.), *Political judgment: structure and process*. University of Michigan Press, Ann Arbor, MI, pp. 91–109.
- Boudreau, C., 2009. Closing the gap: when do cues eliminate differences between sophisticated and unsophisticated citizens? *Journal of Politics* 71 (3), 1–13.
- Bullock, J.G., 2009. Partisan bias and the Bayesian ideal in the study of public opinion. *Journal of Politics* 71, 1109–1124.
- Downs, A., 1957. *An Economic Theory of Democracy*. Harper Row, New York.
- Fazio, R.H., 1995. Attitudes as object-evaluation associations: determinants, consequences, and correlates of attitude accessibility. In: Petty, R.E., Krosnick, J.A. (Eds.), *Attitude strength: antecedents and consequences*. Erlbaum, Mahway, NJ, pp. 247–282.
- Fiorina, M.P., 1990. Information and rationality in elections. In: Ferejohn, J.A., Kuklinski, J.H. (Eds.), *Information and Democratic Processes*. University of Illinois Press, Urbana.
- Fehr, E., Gächter, S., 2002. Altruistic punishment in humans. *Nature* 415, 137–140.
- Festinger, L., 1957. *A theory of cognitive dissonance*. Stanford University Press, Palo Alto, CA.
- Fischbacher, U., 2007. z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics* 10, 171–178.
- Gerber, Green, 1999. Misperceptions about perceptual bias. *Annual Review of Political Science* 2, 189–210.
- Huckfeldt, R., 1983. Social contexts, social networks, and urban neighborhoods: environmental constraints on friendship choice. *The American Journal of Sociology* 89, 651–669.
- Huckfeldt, R., 2001. The social communication of political expertise. *American Journal of Political Science* 45, 425–438.
- Huckfeldt, R., Johnson, P.E., Sprague, J., 2004. *Political Disagreement*. Cambridge University Press, New York.
- Huckfeldt, R., Sprague, J., 1995. *Citizens, politics and social communication: information and influence in an election campaign*. Cambridge University Press, New York.
- Huckfeldt, R., Mendez, J.M., 2008. Moths, flames, and political engagement: managing disagreement within communication networks. *Journal of Politics* 70, 83–96.
- Huckfeldt, R., Sprague, J., Levine, J., 2000. The dynamics of collective deliberation in the 1996 election: campaign effects on accessibility, certainty, and accuracy. *American Political Science Review* 94, 641–651.
- Katz, E., 1957. The two-step flow of communication: an up-to-date report on a hypothesis. *Public Opinion Quarterly* 21, 61–78.
- Katz, E., Lazarsfeld, P.F., 1955. *Personal influence: the part played by people in the flow of mass communications*. Free Press, New York.
- Kearns, M., Judd, S., Tan, J., Wortman, J., 2009. Behavioral experiments on biased voting in networks. *Proceedings of the National Academies of Sciences* 106, 1347–1352.
- Kunda, Z., 1999. *Social cognition: making sense of people*. MIT Press, Cambridge, MA.
- Lazer, D., Rubineau, B., Katz, N., Chetkovich, C., Neblo, M., 2010. The coevolution of networks and political attitudes. *Political Communication* 27, 248–274.
- Levitan, L.C., Visser, P.S., 2009. Social network composition and attitude strength: exploring the dynamics within newly formed social networks. *Journal of Experimental Social Psychology* 45, 1057–1067.
- Lodge, M., Steenburgen, M.R., Brau, S., 1995. The responsive voter: campaign information and the dynamics of candidate evaluation. *American Political Science Review* 89, 309–326.
- Lodge, M., Taber, C., 2000. Three steps toward a theory of motivated political reasoning. In: Lupia, A., McCubbins, M.D., Popkin, S.L. (Eds.), *Elements of Reason: Cognition, Choice, and the Bounds of Rationality*. Cambridge University Press, New York.
- Lupia, A., McCubbins, M.D., 1998. *The democratic dilemma: can citizens learn what they need to know?* Cambridge University Press, New York.
- McKelvey, R., Ordeshook, P., 1986. Information, electoral equilibria, and the democratic ideal. *Journal of Politics* 48, 909–937.
- McPhee, W.N., Smith, R.B., Ferguson, J., 1963. A theory of informal social influence. In: McPhee, W.N. (Ed.), *Formal Theories of Mass Behavior*. Free Press, New York.
- Ostrom, E., Walker, J., Gardner, R., 1992. Covenants with and without a sword: self governance is possible. *American Political Science Review* 86, 404–417.
- Rogers, W., 1993. Regression standard errors in clustered samples. *Stata Technical Bulletin* 13, 19–23.
- Ryan, J.B., 2011. Social networks as a shortcut to correct voting. *American Journal of Political Science* 55, 753–766.
- Verba, S., 1961. *Small Groups and Political Behavior: a Study of Leadership*. Princeton University Press, Princeton.
- Williams, R.L., 2000. A note on robust variance estimation for cluster-correlated data. *Biometrics* 56 (2), 645–646.
- Wolfinger, R.E., Rosenstone, S.J., 1980. *Who votes?* Yale University Press, New Haven.