(SCREEN THIRTEEN) Once again, you are reminded of your previous estimate. If you purchased another piece of private information, it will show up on this screen. Now, it is time to vote. You should vote for the candidate that you think will be closer to your issue position. Your final cash payoff is calculated by adding what is left of your initial endowment to the 50 ECU bonus you receive from the candidate closer to your issue position winning or subtracting the 50 ECU penalty you receive from the other candidate winning. For example, if you had 80 ECUs left from your initial endowment after purchasing private information, and your ideal candidate won, you would end the round with 130 ECUs (80 plus 50). If the other, less ideal, candidate won, you would end the round with 30 ECUs (80 minus 50).

Vote for one of the two candidates and click OK.

(SCREEN FOURTEEN) This is the final screen. The two candidates' positions are revealed as is the outcome of the election. You will also learn the number of ECUs you earned in this period as well as the number of ECUs you have earned up to this point in the experiment.

The experiment will consist of 10 periods like this one. At the end of these 10 periods, you will be asked a couple of questions about the experiment, asked to provide some demographic information, and a couple of questions about your general political leanings. All of your responses are anonymous.

This concludes the demonstration screens. We are now ready to begin the actual experiment. We ask that you follow the rules of the experiment. Anyone who violates the rules may be asked to leave the experiment with only the $5 show up fee. Are there any questions before we start?

10

The complex dynamics of political communication

Robert Huckfeldt, Matthew T. Pietryka, and Jack Reilly

> Each ant lives in its own little world, responding to the other ants in its immediate environment and responding to signals of which it does not know the origin. Why the system works as it does, and as effectively as it does, is a dynamic problem of social and genetic evolution.
>
> Thomas Schelling (1978: 21)

Particularly in the context of complex political processes involving hundreds or thousands or millions of citizens, the whole is typically an unintentional byproduct viewed from the vantage point of the participant. Just as the formation of political beliefs and opinions is not solely due to a cognitive process occurring between the ears of isolated individuals, so too the implications of political communication among citizens is not solely due to an isolated process occurring within self-contained dyads. Not only do the beliefs of individuals depend on what happens within dyads, but the effects of single dyads are contingent on the other dyads within which individuals are simultaneously located (Huckfeldt, Johnson, and Sprague 2004). Moreover, these network effects are not simply cumulative across an individual's range of contacts. To the contrary, the effects are sequential, dynamic, and interdependent. While voters certainly do not resemble Schelling's ants, public opinion in the aggregate is created through complex processes of interaction and communication, located in both space and time, which are at least as complex as those producing the anthill.

This chapter takes a modest step toward understanding an important micro–macro problem in democratic politics (Eulau 1998). In particular, our concern is whether individual levels of political expertise serve to inform the aggregate through the patterns of communication existing among interdependent individuals. We address this problem by extending the analysis of the small group experiments in Chapter 9 to address the consequences of dynamic interdependence for aggregate rather than individual outcomes.

EXPERT CITIZENS AND HIGHER-ORDER COMMUNICATION EFFECTS

Complex patterns of communication and inter-connectedness create analytic opportunities as well as methodological challenges. Opportunities arise to move seamlessly back and forth between aggregates and individuals, guided by the observed patterns of connections between and among the individual members of aggregate populations. At the same time, this potential carries with it a variety of observational challenges. These challenges include endogeneity problems that are endemic to any effort aimed at studying individuals within their ongoing patterns of communication and social interaction – problems that plague any effort aimed at establishing causality in post hoc observational studies of social and political influence. Without experimental control over the flow of communication, it becomes difficult to make assertions that are not hotly contested regarding the effects of expertise on either the flow of communication or its influence, and hence observational studies confront substantial problems with respect to their internal validity.

While post hoc observational studies make it difficult to address endogeneity problems, they also pose problems for studying the higher-order consequences of individual expertise. While inroads have been made in addressing the implications of expertise for political influence and the formation of relationships at the level of dyads, less has been accomplished in addressing the diffusion of expertise through larger populations (but see Nickerson 2008). What are the implications for you if your life partner regularly discusses politics with a knowledgeable person at his or her work place? What are the implications for your life partner that arise due to the coworkers with whom *you* discuss politics?

Moreover, while the social communication of political information carries the potential to create efficiency gains, it also creates the potential for politics and persuasion to be played out in countless social exchanges, with consequences reaching far beyond the immediacy of a dyad (Christakas and Fowler 2009). As we have argued, the social transmission of political information is not an antiseptic exercise in civic betterment. Rather, it is a process characterized by informational asymmetries among participants, as well as frequently passionate advocacy on the part of those who are politically engaged. Thus, it creates the potential for opinion leadership and the social mobilization of bias within the communication process – a process that extends far beyond dyads to generate policy moods (Stimson 1999) and other aggregate consequences.

As a consequence, some individuals are more influential than others, and our goal in this chapter is to assess the higher-order implications of their influence. In particular, we are interested in the extent to which influence reaches beyond dyads. How does political influence extend beyond the immediate range of contacts to penetrate the larger population? Does this penetration serve to amplify the influence of the activists and experts? Here again we confront complex endogeneity problems that are not easily resolved in the absence of an experimental design (Erbring, Goldenberg, and Miller 1980), and hence we extend the analysis of the experiment and experimental results taken from Chapter 9.

A DEGROOT MODEL OF SOCIAL INFLUENCE

The analysis of this chapter employs a DeGroot model (DeGroot 1974; Jackson 2008) to focus on the higher-order consequences of complex communication processes. While we expect the process to be contingent on the preferences and expertise of informants and message recipients, we are less concerned with the direct effects that occur within dyads, and more concerned with the socially mediated effects that arise due to the informants of informants. In the previous chapter we focused on the individual updating process within rounds. In this chapter we employ those results, but our concern turns to the higher-order, longer-term dynamic implications.

The DeGroot model draws on basic theorems regarding Markov chains, where individuals formulate prior beliefs and then update these beliefs on the basis of information taken from other individuals. The updating process is not random, but rather occurs through networks of communication within a larger population. The basic model is

$$p_{t+1} = Tp_t \qquad (10.1)$$

where:

- p_t = is a N×1 column vector, where each of the entries is an individual's belief regarding a particular candidate and p_o is a vector of individual priors. For example, each entry might be the n^{th} individual's belief regarding the position of Candidate A.
- T = a row stochastic matrix, such that T_{ij} is the persuasive weight of the j^{th} individual's belief regarding the candidate at any time period (t) on the i^{th} individual's belief at the subsequent period (t+1). Each row sums to unity, where the main diagonal is the weight that the i^{th} individual places on her own prior belief (at t) in the formulation of her current updated belief (at t+1).

We are especially interested in the T matrix, as well as its long-term dynamic consequences. The focus of this analysis is not on the formation of initial beliefs, but rather on the relative weights that individuals place on their own prior beliefs versus the beliefs of others. In particular, we are concerned with the ways in which the evolution of beliefs in the aggregate depends on the underlying distribution of expertise within the aggregate. And we can readily obtain an estimate of T based on our experimental data from Chapter 9.

TABLE 10.1. *Subject's final judgment regarding Candidate A at each round by their initial (prior) judgment as well as the information conveyed by each of their informants. (Least squares models absent intercepts. Standard errors are adjusted for clustering on subjects.)*

A. All subjects, with no weights for information purchases.

	Coefficient	t-value
Prior judgment	.51	12.61
First message	.17	9.46
Second message	.14	5.36
Third message	.17	5.80
	N =	749 (84 subjects)
	R^2=	.92
	Root MSE=	1.06

B. Subjects who purchased more than 1 piece of information on candidates.

	Coefficient	t-value
Prior judgment	.67	15.12
First message	.10	4.90
Second message	.10	3.85
Third message	.14	5.05
	N =	454 (74 subjects)
	R^2=	.95
	Root MSE=	.86

C. Subjects who purchased less than 2 pieces of information on candidates.

	Coefficient	t-value
Prior judgment	.35	5.95
First message	.23	6.39
Second message	.16	3.85
Third message	.21	3.46
	N =	295 (59 subjects)
	R^2=	.89
	Root MSE=	1.23

In Table 10.1 we address the subjects' final judgments regarding the position of Candidate A at each round.[1] In each part of the table, these judgments are

[1] The results for Candidate B are highly comparable, and the empirical results in Table 10.1 can be compared to those of Table 9.2. We suppress the constant in order to translate the combined effects approximately into a unit interval.

regressed on the subject's prior judgment and each of the three messages that the individual obtained from other subjects. Part A of the table shows the pattern of simple direct effects for all subjects, independently of individual information levels. The coefficients suggest that the effect of the prior is roughly equal to the cumulative effect of the three communicated messages, with comparably sized message effects. In parts B and C of Table 10.1, the model is re-estimated for high- and low-information consumers respectively. We see that the effect of the prior judgment is nearly twice as large among the more informed subjects, with an average message effect that is almost twice as large among the less informed subjects. We explore the aggregate, dynamical implications of these results in the remainder of this paper.

ESTIMATES FOR THE MODEL

The first step is to arrive at estimates for the rows of T – the relative weights that are attached to an individual's own immediately previous judgment, as well as the judgments of others, in arriving at that individual's contemporaneous judgment. The empirical results in Part A of Table 10.1 suggest that, for the subjects as a combined group, approximately 50 percent of current judgments are based on the immediately prior judgments, with the remainder depending approximately equally on the three messages obtained from other participants in the experiment. Hence, the non-zero elements of each row in the T matrix should consist of .5 on the diagonal, with three of the remaining six entries set to .167. (The final message effect is set to .166 in order that the rows sum to 1.)

This raises the obvious question, which three entries? We pursue the objective of considering the long-term implications of the communication choices selected by the participants in two randomly chosen sessions – round 6 of session 5 and round 3 of session 8.[2] These sessions are shown in the directed network graphs of Figure 10.1, where each node is numbered according to the individual's preference, and where the size of the node is indexed on the amount of information purchased by the individual.

These network graphs tell a similar story to the empirical results of Table 9.1. Larger nodes (better-informed participants) attract more requests for information, and higher levels of communication generally occur among individuals with similar preferences. A close inspection of the graphs produces some surprises, and many of these seeming aberrations can be explained on the basis of expertise and preference proximity as competing criteria. This reflects a reality where the choice of informants is complex, and the process is inherently

[2] Each of 12 sessions involved seven subjects and the number of rounds (each of which constituted a separate election) varied from 7 to 10.

stochastic.[3] For the purposes of illustration, the T matrix implied by Part A of Table 10.1 and Part A of Figure 10.1 is shown below.

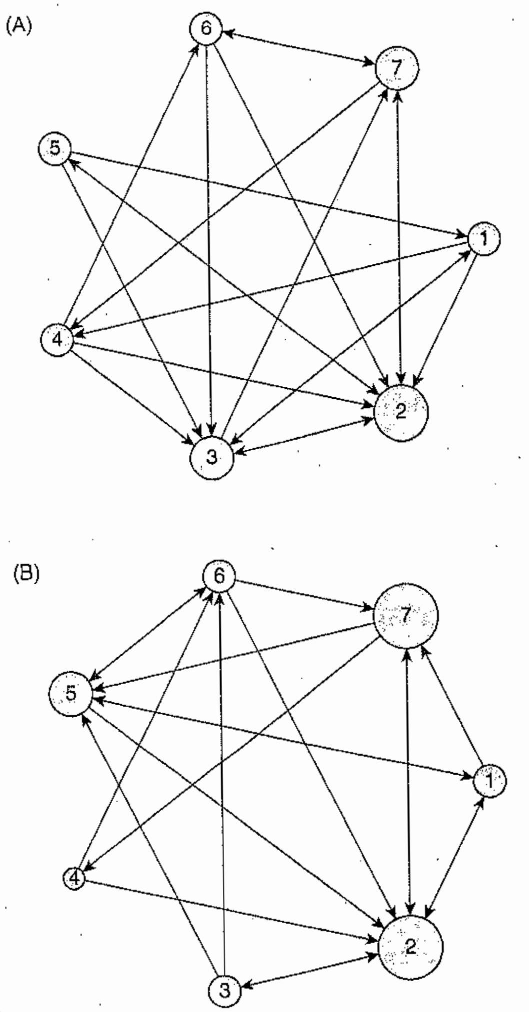

FIGURE 10.1. Directed graphs for randomly chosen rounds. Size of node indexes amount of information purchased. Direction of edge signifies the participant from whom information is being requested.
A. Session 5, Round 6
B. Session 8, Round 3

[3] The stochastic component of participant choices is well illustrated in these two sessions of the experiment. If we regress information purchases on information costs for all the subjects in each round of every session, the R^2 is .23. In contrast, the R^2 for the session in Figure 10.1A is .05 and

$$T = \begin{matrix} .500 & .167 & .167 & .166 & .000 & .000 & .000 \\ .000 & .500 & .167 & .000 & .167 & .000 & .166 \\ .167 & .167 & .500 & .000 & .000 & .000 & .166 \\ .000 & .167 & .167 & .500 & .000 & .166 & .000 \\ .167 & .167 & .166 & .000 & .500 & .000 & .000 \\ .000 & .167 & .167 & .000 & .000 & .500 & .166 \\ .000 & .167 & .000 & .167 & .000 & .166 & .500 \end{matrix} \tag{10.2}$$

Each row is a vector of weights corresponding to the network's contemporaneous, single-period effect on a particular individual's judgment update. Based on the results of Table 10.1A, we set each participant's current judgment as the weighted sum of the immediately prior judgments, with a 50 percent weight on the immediately prior judgment, and the remaining 50 percent partitioned equally among the three informants. Hence, each individual in this specification is influenced directly by her own prior judgment, as well as the judgment of three other participants. Each column corresponds to the contemporaneous, single-period effect of a particular individual's judgment on each of the other individuals' judgments. For example, the fifth column characterizes the very limited short-term effect of the individual who holds preference 5 in Part A of Figure 10.1. Only the individual with preference 2 requests information from the person who holds preference 5.

These short-term direct effects are not, however, simply additive across time. Instead, the information provided by preference holder 5 to preference holder 2 produces indirect effects on all 6 individuals (including preference holder 5) who request information from preference holder 2 at the subsequent time period. Hence, the short-term contemporaneous effects ignore much of the dynamic interdependence underlying communication and persuasion. In this experiment, we do not require individuals to maintain the same contacts across all the experimental sessions, and thus we do not empirically trace effects across the entire session. Rather, our intent is to consider the long-term implications of the short-term contacts that are established by the participants in the experiment, in an effort to consider both the direct and indirect effects of experts and information in political communication.

LONG-TERM DYNAMICS OF POLITICAL COMMUNICATION

The process described in Equation 10.1 is recursive, and hence

the R^2 for the session in Figure 10.1B is .20. In short, while the information cost incentives we have established are clearly related to information purchases, the strategic choices of the participants include a significant idiosyncratic component.

$$
\begin{aligned}
p_1 &= Tp_0 \\
p_2 &= Tp_1 = T^2p_0 \\
&\;\;\vdots \\
p_t &= T^tp_0
\end{aligned}
\qquad (10.3)
$$

This is a particularly helpful formulation because it suggests that each entry in T^t provides the long-term effect of the jth individual in moving the i^{th} individual from p_0 to p_t. If at any value of t, some column of T includes all non-zero elements, it suggests that every individual in the entire population is affected, either directly or indirectly, by every other individual in the entire population, and thus the process leads to convergent beliefs. We are thus assured that the T matrix for round 6 of session 5 leads to convergent beliefs because, in the first order matrix, the column for preference holder 5 contains all non-zero elements, and T raised to successively higher powers converges on the following set of identical row vectors.

$$
T^* = \begin{matrix}
.096 & .250 & .205 & .093 & .084 & .091 & .181 \\
.096 & .250 & .205 & .093 & .084 & .091 & .181 \\
.096 & .250 & .205 & .093 & .084 & .091 & .181 \\
.096 & .250 & .205 & .093 & .084 & .091 & .181 \\
.096 & .250 & .205 & .093 & .084 & .091 & .181 \\
.096 & .250 & .205 & .093 & .084 & .091 & .181 \\
.096 & .250 & .205 & .093 & .084 & .091 & .181
\end{matrix}
\qquad (10.4)
$$

Hence we can identify an equilibrium vector for participants' judgments, which is simply the judgmental priors of the participants weighted by T^*.

$$p^* = T^*p_0 \qquad (10.5)$$

Assume for the moment that p_0 provides the participants' initial estimates regarding the position of Candidate B, specified as (2,2,2,3,4,4,4). This leads to an equilibrium vector of (2.8, 2.8, 2.8, 2.8, 2.8, 2.8, 2.8). If we reversed the order of priors in p_0, p^* would become an equilibrium vector (3.2, 3.2, 3.2, 3.2, 3.2, 3.2, 3.2). In short, these alternative outcomes reflect the greater influence of those individuals with preferences 1 through 3 relative to preferences 5 through 7.

A row of T^* becomes the unit eigenvector for T,[4] and it provides a relative measure of each participant's influence, capturing both their direct and indirect

[4] A row vector in T^* provides the unit eigenvector of T – the row vector that, when multiplied times T, returns the same row vector, or tT=t.

TABLE 10.2. *Unit eigenvectors for experimental periods. Expert effects in bold italics.*

	Participant with preference:							
	1	2	3	4	5	6	7	Σ expert
A. Baseline condition.								
Group 5, Period 6:	(.096	***.250***	***.205***	.093	.084	.091	***.181***)	.64
Group 8, Period 3:	(.141	***.250***	.084	.057	***.173***	.123	***.171***)	.59
B. With information weights.								
Group 5, Period 6:	(.059	***.305***	***.249***	.057	.052	.056	***.222***)	.78
Group 8, Period 3:	(.090	***.312***	.053	.036	***.217***	.078	***.214***)	.74

effects within the network. By comparison to Figure 10.1A, we see that the individuals who purchased the most information become the opinion leaders in the process. Indeed, the participants in columns 2, 3, and 7 combine for 64 percent of the total network influence.

The network in Part B of Figure 10.1 produces a similar outcome. Part A of Table 10.2 shows the unit eigenvectors for both networks. In each instance, the individuals who purchased the most information demonstrate the strongest relative effects. Indeed, as shown in part A of Table 10.2, the three highest consumers of information in the two randomly selected experimental rounds account for 64 and 59 percent of opinion leadership, respectively.

These networks are inherently stochastic. Individuals need not purchase information, and even individuals who can obtain it for free do not necessarily obtain the full amount that is available. The relationship between information costs and preference is the same in both experimental periods, yet we see variation between the rounds in information purchases. No one in Part A obtained four pieces of private information, but two individuals in Part B purchased the maximum. In short, information costs and preference proximity generate important effects on network structure, but these effects are certainly not deterministic, and we see pronounced differences in networks between the two experimental periods.

Finally, it is important to recognize that the empirical model of opinion leadership in Part A of Table 10.1 is wholly due to the specification of network selection. That is, we are assuming network effects that are wholly mediated by an individual's choice of discussion partners, without any effects due to the inherent effects of information abundance and scarcity on the processing of either private information or socially communicated information. We turn to the consequences of information and expertise for the behavior of individuals, both with respect to the confidence they place in their own priors, as well as the extent to which they update their priors based on socially communicated information.

INFORMATION, EXPERTISE, AND OPINION LEADERSHIP

As Parts B and C of Table 10.1 suggest, information has important effects on the extent to which individuals depend on their own priors versus the extent to which they depend on messages received from other individuals. That is, individuals who purchase more information reveal more confidence in their own prior judgments, as well as relatively less confidence in the judgments of others. This moves us beyond a strictly sociological view of the problem based on the structure of the relationships among participants, introducing a psychological perspective regarding the cognitive processing of new information, as well as a decision-making process whereby individuals update their own preconceived judgments in a highly uncertain environment with information provided by others.

Unfortunately, major advances in the cognitive processing of information are rarely considered among network scientists, and major advances in network science appear to be unknown among students of cognition. It is as if network scientists ignore the nodes, and cognitive scientists ignore the edges![5] The reality is that, as Parts B and C of Table 10.1 show, it is not simply the existence of the communication pathways among actors that is important. Indeed, the relative efficacy of those pathways depends on the commitment of individual actors to their own pre-existent beliefs, as well as their relative openness to communication on the part of the individuals (nodes) populating the pathways. These results reinforce the work of Lodge and Taber (2000) – those who know the most are the least willing to be moved by new information.

We build on these results by constructing T matrices that take account of these contingent effects for both network graphs in Figure 10.1. The matrix for Figure 10.1A is shown here:

$$T = \begin{matrix} .350 & .217 & .217 & .216 & .000 & .000 & .000 \\ .000 & .670 & .110 & .000 & .110 & .000 & .110 \\ .110 & .110 & .670 & .000 & .000 & .000 & .110 \\ .000 & .217 & .217 & .350 & .000 & .216 & .000 \\ .217 & .217 & .216 & .000 & .350 & .000 & .000 \\ .000 & .217 & .217 & .000 & .000 & .350 & .216 \\ .000 & .110 & .000 & .110 & .000 & .110 & .670 \end{matrix} \qquad (10.6)$$

Reflecting the results of Parts B and C of Table 10.1, the better-informed rely more heavily on their priors and less heavily on social communication than the lesser-informed. This produces, in turn, the unit eigenvector in Part B of Table 10.2, which displays an enhanced level of influence for the informed relative to the uninformed.

[5] For particularly notable exceptions see Levitan and Visser (2009) and Lazer et al. (2010).

In summary, opinion leadership is both a sociological as well as a psychological phenomenon. Not only is the influence of opinion leaders related to their centrality within communication networks and the frequency with which they engage in political communication, but it is also due to the resilience and durability of their judgments. Experts are resistant to persuasion and committed to their own prior judgments, thereby providing a persuasive advantage in the collective deliberations of democratic discussion.

THE DECISIVE EFFECTS OF SLOWLY DECAYING PRIORS

Finally, limits on cognition encourage us to take the mechanics of memory seriously in the analysis of political communication. Working memory is dramatically limited in its capacity, and objects in working memory can be lost after they are passed to long-term memory. Hence, memory decay plays a role in the duration of even the strongest beliefs and judgments.

Analyses of these same experimental results support even short-term effects on memory decay (see Chapter 9) – that is, participants in the experiment update their prior judgments three times during a round, and during that short period of time we see a rate of decay in the priors that is especially precipitous among the least informed. At the same time, even the priors of the most informed show the short-term consequences of memory decay.

Our goal is to consider the implications that arise due to differential rates of decay among experts and non-experts, and we modify the general model accordingly (see Friedkin and Johnsen 1990; Jackson 2008). Suppose that an individual's prior judgment competes directly with the updating process, but that the importance of the prior declines in time as a consequence of memory decay. We incorporate this idea into a revised model,

$$p_{t+1} = D^{t+1}p_0 + (I - D^{t+1})Tp_t \qquad (10.7)$$

where the T matrix is taken from (10.6); D is a diagonal matrix with the rates (defined on a 0,1 interval) at which individuals' prior judgments survive in a single period of time; I is the identity matrix; and I-D is a diagonal matrix reflecting the complement of D – the rates at which individuals base their judgments on the messages received from informants ***as well as*** their own immediately previous judgments.

Hence, the importance of an individual's prior declines both in time and across individuals. For the purposes of illustration, we set the rate of decay to .2 for individuals who purchased 2 or more pieces of information, and to .6 for individuals who purchased 0 or 1 piece of information. (These rates of decay are compatible with the earlier analyses in Chapter 9 showing rates of decay structured by private investments in information.)

First, the impact of memory decay declines over time in this formulation – that is, D^{t+k} converges to zero as k increases, but it converges more rapidly among the least informed. Second, and in a similar fashion, $I-D^{t+k}$ converges on I (the identity matrix), as k increases. And hence, in the long run, the effect of memory decay disappears and we are left with the process described in the basic model (see Equation 10.1). The end result is the same unit eigenvector that is displayed in Table 10.2B for Group 5, Session 6. The difference is that the system takes much longer to converge on the same equilibrium vector of shared judgments, and ***during that slow path to equilibrium, the judgments of the opinion leaders are more influential.***

The question that arises is whether the long term makes much difference in the deliberations of democratic societies, and the answer is a resounding "sometimes"! Many issues play out on a short timescale, and in these instances opinion leaders are likely to be particularly influential because their priors decay so slowly. Other issues are of longer duration, and we should expect an inevitable convergence toward a long-term equilibrium which serves to diminish the role of enduring priors even among opinion leaders. Indeed, some issues are initially controversial, but eventually become settled matters after long periods of public discussion. Universal suffrage is one example, and the existence of global warming is likely to be another. This does not mean that everyone ends up holding the same opinion, but in these cases divergence from the conventional view become notable because they are so rare and idionsyncratic.

INCORPORATING THE MODEL OF MEMORY DECAY

While memory constraints certainly operate on prior judgments, they also operate on the updates to these priors, as well as the messages that are communicated by others. Memory decay can be portrayed at the individual level 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) incoming information communicated by other subjects. For the reader's convenience, we briefly recapitulate the model developed in Chapter 9.

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_o, where w is defined as (1–rate of decay) and at the n^{th} update its effect is thus w^nP_o.

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 updared judgments are subject to decay, the subject is responding to an ongoing stream of social information communicated by other subjects. This incoming information is incorporated within the update, and thus its effect decays as the updated judgment decays.

Hence, the current judgment arises due to the persistence of the immediately preceding judgment update, the rate of decay in an initial prior judgment, and the effect of contemporaneous social information.

$$J_t = \alpha J_{t-1} + w^t P_o + eI_t \text{ , or} \tag{10.8}$$

where α = the memory or survival of the previous judgment ($0<\alpha<1$); P_o = the initial or prior judgment based on privately purchased information; w^t = the effect of the prior at t ($0<w<1$); I_t = socially supplied information received at t; and e = the educative impact of the new social information.

By employing recursion to push the model beyond the reach of our experimental data, we take the equation to its limiting behavior. For n sufficiently large,

$$J_n = (w^n P_o - \alpha^n w^o P_o)/(1 - \alpha/w) + eI_n + \alpha eI_{n-1} + \ldots + \alpha^{n-1} eI_1. \tag{10.9}$$

Assuming that both α and w are bounded hy 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.

The crucial issue is how rapidly the memory of this behavioral system decays. 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 wherehy the prior is modified by new information, α also provides an index of 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.

OBTAINING ESTIMATES FOR THE DEGROOT MODEL

The three model parameters are re-estimated for low- and high-information subjects relative to both candidates in Table 10.3, based on the procedures outlined in Chapter 9. Part A of Table 10.3 displays the results of estimating the model in Equation 10.8 for high-information subjects, defined as subjects who purchased more than 1 piece of information. Part B shows the results for low-information

TABLE 10.3. *Final judgment regarding candidate positions by initial prior judgment, immediately previous judgment, and previous (third) message received from other participants, for high-information and low-information subjects.*
A. High-information subjects who purchased more than 1 piece of information.

	Candidate A		Candidate B	
	Coefficient	T-value	Coefficient	T-value
Prior	.19	2.85	.28	4.07
Previous judgment	.65	8.88	.60	7.27
Previous message	.12	5.32	.10	5.09
Constant	.12	1.22	.12	1.05
N=	454		454	
Subjects=	74		74	
R2=	.79		.83	
Root MSE=	.67		.62	
	parameters			
	w=.57		w=.65	
	α=.65		α=.60	
	e=.12		e=.10	

B. Low-information subjects who purchased less than 2 pieces of information.

	Candidate A		Candidate B	
	Coefficient	T-value	Coefficient	T-value
Prior	.02	.59	.06	4.07
Previous judgmenr	.67	9.01	.64	7.27
Previous message	.18	4.15	.20	5.09
Constant	.33	1.98	.41	1.66
N=	295		295	
Subjects=	59		59	
R2=	.55		.53	
Root MSE=	.99		1.07	
	Paramerers			
	w=.27		w=.39	
	α=.67		α=.64	
	e=.18		e=.20	

subjects who purchased less than 2 pieces of information. In each case, the final updated judgment (J_3) is regressed on the immediately preceding updated judgment (J_2), the initial prior judgment (P_o), and the immediately preceding (third) piece of communicated information (I_3).

The estimated model parameters are consistent with the Chapter 9 results, showing that the effect of the initial (prior) judgment is dramatically dependent on the amount of information purchased by a subject – the effect of the prior persists only among those participants who invest in private information. The table also shows a substantial effect due to the immediately preceding update that is comparable among high- and low-information individuals. Finally, we see a substantial effect due to the contemporaneous message that is attenuated by the amount of private information purchased by the participant.

Returning to the modified DeGroot model in Equation 10.7, not only can we specify the rates of memory decay in the priors, but we can also take account of memory decay with respect to the social communication process that is captured in the T matrix. First, the effect of the immediately preceding judgment is α, and the effect of the incoming information is e. All the model parameters are represented in the T matrix as summing to unity across the rows. This involves adjusting the parameter magnitudes so that their effects relative to one another are maintained. Hence, at each social iteration, the normed effects for α and e are set at $\alpha/(\alpha+e)$ and $e/(\alpha+e)$.

As the model in Equation 10.9 suggests, earlier messages are subject to decay. The most recent message effect is "*e*," and earlier messages are discounted by raising α^{n-1} to successively higher powers. This occurs automatically, as the updated judgments which incorporate the new information are adjusted by α at each iteration.

The model assumes that each individual has three informants, and that each individual cycles through the informants in the order in which they were initially chosen, receiving and responding to each message in the order in which it is received. Hence, the first informant sends messages at t = 1, 4, 7, etc. The second informant sends messages at t = 2, 5, 8, etc. And the third informant sends messages at t = 3, 6, 9, etc. Correspondingly, we construct three T matrices, corresponding to the subjects' first, second, and third choices of informants.

Hence the T_1 matrix becomes:

$$\begin{array}{llccccccc}
 & & \multicolumn{7}{c}{\text{Informant Preference}} \\
 & & 1 & 2 & 3 & 4 & 5 & 6 & 7 \\
 & 1 & .79 & .00 & .21 & .00 & .00 & .00 & .00 \\
 & 2 & .00 & .84 & .00 & .00 & .00 & .00 & .16 \\
\text{Requestor} & 3 & .00 & .16 & .84 & .00 & .00 & .00 & .00 \\
\text{Preference} & 4 & .00 & .00 & .21 & .79 & .00 & .00 & .00 \\
 & 5 & .00 & .00 & .21 & .00 & .79 & .00 & .00 \\
 & 6 & .00 & .21 & .00 & .00 & .00 & .79 & .00 \\
 & 7 & .00 & .00 & .00 & .00 & .00 & .16 & .84
\end{array} \qquad (10.10)$$

The T_2 matrix becomes:

	Informant Preference							
		1	2	3	4	5	6	7
Requestor Preference	1	.79	.00	.00	.21	.00	.00	.00
	2	.00	.84	.16	.00	.00	.00	.00
	3	.00	.00	.84	.00	.00	.00	.16
	4	.00	.21	.00	.79	.00	.00	.00
	5	.00	.21	.00	.00	.79	.00	.00
	6	.00	.00	.00	.00	.00	.79	.21
	7	.00	.16	.00	.00	.00	.00	.84

(10.11)

And the T_3 matrix becomes:

	Informant Preference							
		1	2	3	4	5	6	7
Requestor Preference	1	.79	.21	.00	.00	.00	.00	.00
	2	.00	.84	.00	.00	.16	.00	.00
	3	.16	.00	.84	.00	.00	.00	.00
	4	.00	.00	.00	.79	.00	.21	.00
	5	.21	.00	.00	.00	.79	.00	.00
	6	.00	.00	.21	.00	.00	.79	.00
	7	.00	.00	.00	.16	.00	.00	.84

(10.12)

where the rows and columns for well-informed subjects are shown in bold parentheses.

In a similar fashion, we can estimate the weights that subjects place on their own priors (w), and on this basis construct the D matrix in Equation 10.7. As we will see, this dramatic difference carries important consequences for the social dynamic, and it produces the following D matrix.

$$D = \begin{matrix} .27 & .00 & .00 & .00 & .00 & .00 & .00 \\ .00 & .57 & .00 & .00 & .00 & .00 & .00 \\ .00 & .00 & .57 & .00 & .00 & .00 & .00 \\ .00 & .00 & .00 & .27 & .00 & .00 & .00 \\ .00 & .00 & .00 & .00 & .27 & .00 & .00 \\ .00 & .00 & .00 & .00 & .00 & .27 & .00 \\ .00 & .00 & .00 & .00 & .00 & .00 & .57 \end{matrix} \qquad (10.13)$$

Based on these D and T matrices, as well as Equation 10.7, we can iteratively estimate convergence paths for the individual requestors. Figure 10.2a displays

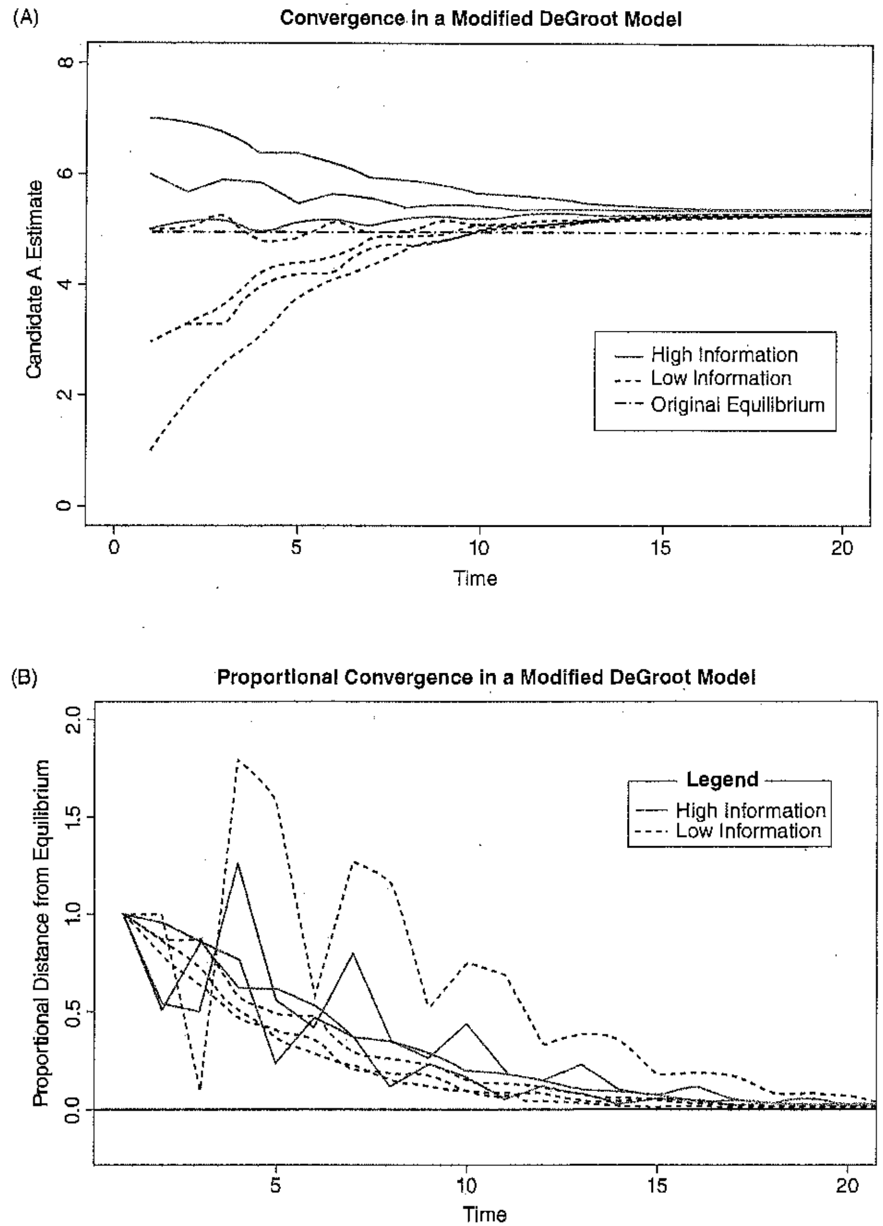

FIGURE 10.2. Convergence to equilibrium.
A. Estimated convergence to equilibrium among subjects by expertise.
B. Estimated proportional convergence to equilibrium among subjects by expertise.

the modified DeGroot estimates of the requestors' judgments regarding candidate positions across time. Highly informed individuals (2, 3, and 7) are shown with solid lines, whereas low-information individuals are shown with dashed lines. A few things are particularly notable about this graph. The highly informed individuals, most of whom begin estimating Candidate A's position at the high end of the scale, are more influential to the creation of the final equilibrium than low-information individuals. The three high-information subjects estimated Candidate A's position as a 6 on average, while the low-information subjects averaged a 3. However, the final equilibrium value estimated by our modified DeGroot model (5.30) is far closer to the estimates of the high-information, rather than the low-information, voters.

The influence of the experts arises as a consequence of two factors. First, high-information voters have a ***much stronger attachment to their priors*** than low-information voters. Second, high-information voters are ***asked to provide information more often*** than low-information voters, so their viewpoints are more influential throughout the network. As can be seen in Figure 10.2a, high-information voters also tend to be slower to converge to equilibrium – especially early on – as their attachment to their prior makes them much more resistant to change.

Figure 10.2b displays the convergence toward equilibrium as a proportion of the distance between the subject's initial prior judgment and the final equilibrium. By definition, every subject starts at "1," and subsequent values greater than 1 indicate that the subject has diverged (moved farther away) from the eventual equilibrium.[6] As before, solid gray lines indicate high-information individuals, whereas dashed black lines indicate low-information individuals. Equilibrium is represented by a solid horizontal line at zero. Hence, the criterion variable standardizes a subject's distance from equilibrium at any point in time relative to the initial distance from equilibrium at the beginning of the process. This serves to enhance the observed magnitude of change among those individuals who begin the process near to the ultimate equilibrium.

In this graph we see that most individuals behave roughly as expected: three low-information individuals, relatively unattached to their priors, converge to equilibrium more rapidly on average than the high-information subjects. The high-information subjects also converge to equilibrium, but the slow decay in their priors delays the convergence.

The seeming exception is the low-information individual whose behavior is characterized by a few wild swings before she or he also begins to converge. This path to equilibrium is explained by the particular patterns of interaction within the communication process – the subject requested information from particular

[6] In theory, values under zero indicate that the subject, after having prematurely converged rapidly upon equilibrium, has overshot and diverged away from equilibrium to the other side. However, there are no cases of this in our currently selected round.

individuals who provided initially divergent signals, pushing the subject farther away from the ultimate equilibrium. (While the swings are not as pronounced, substantial shifts can also be seen in two of the high-information subjects.) While this initial advice was eventually attenuated, the impact of that information persisted due to the subject's reliance on his or her own immediately prior judgments.

Hence, Figure 10.2 serves to illustrate the noisy nature of the communication process. While strong pressures toward equilibrium tend to filter and dampen aberrant messages, particular time paths are highly dependent on particular communication patterns and events, as well as the order in which information is received, leading to highly diverse and variable dynamics across individuals and groups.

IMPLICATIONS AND CONCLUSIONS

Several features of this chapter's argument are particularly important. Perhaps most crucially, the model supports a compelling dynamic analysis regarding the role played by opinion leaders. The role of opinion leaders – the experts in this analysis – is enhanced by two immediate factors. First, they rely more heavily on their own priors and immediately past judgments in responding to socially communicated information, and they are secondarily affected by socially communicated information.

The dynamic implications are quite important. Opinion leaders have sticky priors and past judgments – their prior beliefs persist even under the onslaught of new socially communicated information. This means that opinion leaders are much less likely to dramatically modify their own opinions due to changing opinion distributions in the aggregate. Their staying power, in turn, serves as an anchor on changes in public opinion. Rather than moving toward new opinions, they tend to pull the movement of public opinion back toward their own beliefs.

The persistence of expert priors carries several important implications. Most particularly, the role of experts does not depend on their own loquacious arguments and compelling analyses of public affairs. *Their influence is as much due to their unwillingness to move as it is their ability to encourage movement among others.*

This result helps to make sense of the otherwise puzzling cross-sectional analyses of opinion leadership. These results tend to show weak or non-existent expertise effects on the levels of persuasiveness within dyads. Thus, in Chapter 3, we addressed survey respondents who were very likely to recognize expertise among their discussion partners, but were not necessarily more responsive to the arguments of these experts in comparison to the arguments of their less expert associates. By moving to a dynamic experimental analysis, we are able to address the effect of new socially communicated information within the context of a dynamic process where the participants hold priors, based on varying levels of information and commitment, and are willing to update their opinions accordingly. Such an analysis depends on a highly dynamic treatment of opinion leadership and the influence of opinion leaders.

Finally, the DeGroot model produces a single equilibrium that serves as an attractor for all of the participants. Lacking new information from the surrounding environment, the analysis suggests that a dynamic process occurs in which consensus is reached and that opinions converge toward a particular level. This is both encouraging and discouraging. On the one hand, it conforms to what we have come to see as the news cycle, where (1) new information is communicated through print and electronic media, (2) public responses are initially at disequilibrium, (3) opinions in the population crystallize and stabilize toward a new consensus (see McPhee 1963).

At the same time, this feature of the model poses an inherent limitation – not only is a population equilibrium reached, but this equilibrium consists of a single belief. That is, the model predicts that all groups tied together by direct and indirect ties will inevitably converge toward a consensual equilibrium of shared opinions. At the same time, the empirical record demonstrates the survival of heterogeneous opinions within self-contained populations. Not only do the friends of your friends hold views with which you disagree, but you probably have at least a few friends with whom you are not in perfect harmony. In short, portraying the dynamic that yields a stable, heterogeneous equilibrium is an ongoing challenge in the study of communication networks that reveals the need for non-linear models (Huckfeldt, Johnson, and Sprague 2004) and points toward the inherent complexity of non-linear interdependence (see May's 1976 observation at the beginning of Chapter 9).

Once again, we see the role played by the division of labor in the communication of political information (Berelson et al. 1954). This chapter suggests that the influence of opinion leaders lies in their own unwillingness to change their beliefs. The experts and activists in our midst, correctly or incorrectly, carry with them the courage of their own convictions. The resulting patterns of communication sometimes produce an electorate that makes surprisingly expert choices – at least relative to the low mean levels of political awareness among individuals within the electorate (Converse 1964; Zaller 1992; Delli Carpini and Keeter 1996; Page and Shapiro 1992; Erikson, MacKuen, and Stimson 2002). Alternatively, we have no guarantees that the experts and the activists are always right. Indeed, the experts and activists reading these words would have a difficult time arriving at their own consensus!

When people share their opinions, they are not only learning from one another, but are also persuading one another. The accumulated record, based on surveys and experiments, suggests that the process tends to be driven by activists and experts. That is, the process is skewed in favor of politically engaged participants with more information – the same individuals identified by Lodge and Taber (2000) as being most opinionated and most likely to demonstrate motivated reasoning. Moreover, because the experts are often activists within their own closely held networks of communication – that is, they are not dispassionately neutral – they have the potential to mislead as well as to inform. The information they communicate is typically biased, reflecting the interests of the informant, and thus we cannot assume that all crowds are "wise" crowds (Surowiecki 2005). Even when it comes to "facts" as opposed to "interpretation," not only information diffuses through communication networks, but misinformation as well.

In short, political communication among citizens is not simply an exercise in civic enlightenment – it is instead an inherently political process that plays a central role in democratic politics. Hence, the process carries no guarantees of producing wise or enlightened outcomes. Rather, such a result depends on the continued effort and vigilance of those engaged experts and activists who value democratic outcomes.