Information Sharing in Small Group Deliberation*

Charles D. Myers

Princeton University

April 27, 2009

Abstract

Deliberation is supposed to lead to better, more legitimate public policy by increasing public knowledge. Theorists pin these claims on the way deliberation promotes information sharing, exposing citizens to new facts and perspectives. But will people share information in deliberation? In this paper I combine game-theoretic and psychological theories about information sharing to create a full theoretical account of when information will and will not be shared in deliberation. I test these theories using an experimental design that combines the methodologies of social psychology and experimental economics. Initial results suggest that information sharing in deliberation is rarely perfect, and is

*Thanks to Tali Mendelberg, Adam Meirowitz, Martin Gilens, Chris Achen, Chris Karpowitz, Shawn Rosenberg, Chuck Myers, Dave Glick, Dustin Tingley, Ben Lauderdale and Kristen Walcott for helpful feedback, discussion, and comments on this project. Financial support was provided by the Mamdouha S. Bobst Center for Peace and Justice. All errors, be they of omission, commission, or typography, are of course my own.

biased in predictable ways that should be taken into account when designing deliberative institutions.

> *A well constituted deliberative forum provides an opportunity for advancing both individual and collective understanding. Participants can learn from each other, come to recognize their individual and collective misapprehensions, and develop new views and policy that can withstand critical scrutiny.* (Gutmann and Thompson 2004, pg. 12).

> *New information is imparted because 1) no single individual can anticipate and forsee all of the variety of perspectives through which matters of ethics and politics would be perceived by different individuals; and 2) no single individual can posses all of the information deemed relevant to a certain decision affecting all. Deliberation is a procedure for being informed.* (Benhabib 1996, pg. 71).

Deliberative theories of democracy hold that decisions are only legitimate when they are made in a free and open deliberation where reasons are offered to and by those affected by government policy (Gutmann and Thompson 1996, Cohen 1989). Beyond offering legitimacy, these normative theories claim that deliberative democracy can lead to more just public policy by increasing public knowledge about the moral and practical issues at stake and making citizens aware of the perspectives of socially disadvantaged groups. Democratic theorists pin these claims on the way deliberation promotes information sharing, exposing citizens to facts, opinions, and perspectives that they had not previously considered.[1]

[1] I use the term information broadly. While it includes factual pieces of knowledge, it may also include arguments, new perspectives, or any other statement that might change the preferred outcome of some member of the deliberative body.

But will people share information in deliberation? If they share information, will others treat it seriously? Normative theories of deliberative democracy contain a moral imperative for people in deliberative institutions to make public their reasons for supporting a particular policy and seriously consider the reasons offered by others. However, little empirical work has examined the willingness of actual deliberators to meet these standards. As society begins to create structures that institutionalize the deliberative ideal, we are compelled to take this question more seriously.

This paper presents the initial results of an experiment that explores the causes of incomplete information sharing in deliberation, as well as the consequences for the quality and legitimacy of decisions produced by deliberative institutions. I combine a game-theoretic model of strategic information transmission with psychological theories about information sharing in small group discussion to create a full theoretical account of when information will, and will not, be shared in deliberation. I test the empirical implications of this theory in an innovative experimental design that combines the methodologies of social psychology and experimental economics. Initial results suggest that information sharing in deliberation is rarely perfect, and is biased in predictable ways that should be taken into account when designing deliberative institutions

This paper will proceed in the following manner. Section one reviews the psychological and game-theoretic literatures related to information sharing in group discussion. Section two develops a simple model of information sharing that generates hypotheses that are testable in the lab. Section three describes the experimental design, and section four discusses some initial results from the experiment.

1 Related Literatures

There are two scholarly literatures that suggest that information sharing in deliberation cannot be taken for granted. The first is the game theoretic literature on strategic communication in games of incomplete information. The second literature is the psychological literature on information sharing in group discussion. These literatures suggest that information that is not widely shared before deliberation begins, information like the technical knowledge of experts and the perspectives of minority groups, will have less of an impact on the outcome of deliberation than information that is widely known before deliberation.

Psychologists have long been interested in the way people share information in group decision making contexts. The key question in this literature is how the pre-discussion distribution of a piece of information, whether it is known by one or several group members, affects the use of that information in group decision making. Under ideal circumstances, whether a piece of information was widely or narrowly known before discussion begins would not alter the impact that piece of information had on the final decision. However, groups tend to spend more time discussing information that is known by all discussants before the deliberation begins, and less time on information that is known by only a few group members. Information that is widely known before discussion also has greater influence over the final group decision.

This so-called "common knowledge effect" was first noted by Stasser and Titus (1985). Gigone and Hastie (1993) show that the greater effect of pre-discussion "common knowledge" is a result of more discussion of information that is widely shared before discussion and the fact that this information is incorporated into a greater number of deliberators' pre-discussion preferences. A number of studies vary other factors of the decision making task, including group size (Cruz, Boster and Rodrìguez 1997), whether the task has a clear correct answer (Stasser and Stewert 1992), and the time dedicated to discussion (Parks and Rebecca 1995), but still find a tendency to focus on information that is shared before discussion begins. This common knowledge effect is troubling given the confidence that normative theorists place in discussion as a way to reach just outcomes. If the results of a deliberation are dependant on the distribution of information before deliberation begins, minorities and other groups whose experiences are not well known to the general public will be less able to influence deliberative outcomes.

Two factors have been shown to eliminate the common knowledge effect. The first, and most obvious, solution is to make all information available to all discussants before discussion begins. The second solution is to assign expert roles to discussants with unshared information. When participants with "uncommon" information are identified, this information has the same impact on the group decision as commonly known information (Stasser, Vaughan and Stewart 2000). In my experimental design I rely on both of these factors to manipulate the existence of the "common knowledge effect."

In addition to psychological work on the common knowledge effect, the game-theoretic literature on strategic information transmission predicts incomplete information sharing, though for different reasons. Here, the degree of information sharing is dependant on the degree to which deliberators share the same preferences over outcomes. Beginning with Crawford and Sobel (1982), these models demonstrate that when all players have identical interests costless communication can perfectly convey information in equilibrium. However, if interests are not identical communication is a strategic interaction. Since communication affects each players' payoff indirectly through the actions of the other players, there is an incentive for a deliberator with unique knowledge of a piece of information to either shade the truth or outright lie. Given this incentive, other deliberators are rightly suspicious of all statements, making completely honest communication impossible. Other models expand this insight into a deliberative setting (Austen-Smith 1990, Meirowitz 2006) where multiple players have different private information, or where multiple actors have access to the same private information (Krishna and Morgan 2001), and again demonstrate that communication depends on similarity of interests. Several experimental studies find support for this basic result (Dickhaut, McCabe and Mukherji 1995, Cai and Wang 2006) though they find that participants communicate more information than predicted. However, these studies generally involve highly stylized decisions (for example, guessing randomly generated numbers), and none involve face-to-face communication.

Both of these literatures predict less than full information sharing in deliberation, but for different reasons. Without carefully isolating the causal mechanisms underlying these different reasons it is difficult to tell which is causing a lack of information sharing. Unfortunately, the psychological literature begins with the assumption that all participants share the same preference. Experiments that make this assumption can not tell us the effect of preference heterogeneity. They also run the risk of attributing a lack of information sharing caused by unobserved preference heterogeneity to the common knowledge effect.

Consider, for example, Gigone and Hastie's 1993 study of information sharing in small groups that were told to give grades to students in an introduction to psychology class. Gigone and Hastie implicitly assume that all group members share the same preference: giving the student the correct grade. It is possible, however, that some wished to give generally higher grades (perhaps out of a sense of altruism) while others wished to give generally lower grades (perhaps out of a concern for grade inflation). A participant that favored giving an average grade of B in the class and had unique knowledge of a particular student's poor attendance record might have an incentive to withhold this information if he thought others wanted to give an average grade of C. Since these preferences are not observed, this lack of information sharing would be incorrectly attributed to the common knowledge effect.

Experimental evidence on the effect of preferences on information sharing is clearer than psychological studies of the common knowledge effect largely because these experiments do not include realistic discussion. However, this heightened experimental control is at the cost of external validity. We can not be sure that this effect will persist in more realistic discussion situations (where, for instance, norms against lying might create more honest communication), or compare the magnitude of this effect to the magnitude of the common knowledge effect. By combining these two approaches and carefully controlling both causal mechanisms, the experiment I describe below will fill an important gap in the literature.

2 A Model of Information Sharing in Deliberation

When do strategic incentives make it hard for information to be transmitted in deliberation? Game theorists answer this sort of question by looking for equilibrium in cheap-talk games of incomplete information, or signaling games. Such games usually begin with nature randomly selecting a state of the world that only one player is aware of. Signaling games generally have two types of equilibrium, pooling and separating.[2] In a pooling equilibrium the knowledgeable player takes the same action regardless of the state of the world nature selects, while in a separating equilibrium the player takes different actions depending on the state of the world. We can think of deliberation with private information as a signalling game in which only one player, who I will refer to as the "expert," has more information about nature's action. If a separating equilibrium exists, the expert sends a signal that accurately communicates this information to the non-expert audience, the ideal outcome for deliberation. However, if only a pooling equilibrium exists the expert can not convey any information to the non-expert.

In this section I will describe a fairly simple signaling game that captures the essence of the expert - non-expert interaction. By removing some of the more complicated elements of other models of deliberation while retaining those that are most important to approximating the conditions of deliberation, the model generates empirically testable predictions. In particular, it focuses on generating predictions that are testable alongside the predictions of the common knowledge effect literature. To this end, I incorporate two features not commonly seen in models of communication: uncertain information, and partially state-of-the-world dependant preferences. Few policy debates involve perfect information about the consequences of different decisions. My model captures this by making all information indicate the true state of the world with a certain probability. Similarly, most policy debates involve some elements on which the parties fundamentally disagree, and other elements where rec-

[2] In more complicated games partially separating equilibrium may exist, but not in the game described here.

onciliation is possible. My model captures this by giving each actor a non-negotiable preference for one option but a willingness to compromise if that option is shown to be clearly inferior. In addition, I assume that deliberators only play weakly undominated strategies when voting and that the preferences of all voters are public knowledge.

Consider a committee of three that must decide between two policy options, x and y. The final decision is known as z and is made by majority rule voting. Each player has a preferred policy choice that defines their type θ. These types are public information. Without loss of generality, assume that player 1 is type x, and for convenience arrange the plays so that a $\theta = y$ player is player three if there is a $\theta = y$ player. In this way, player two is always the deciding vote in any situation where types $\theta = x$ and $\theta = y$ disagree.

Each player derives some utility from selecting their preferred policy and some utility from selecting the policy that most closely matches the unknown state of the world ω.[3] Specifically, each players' utility function is:

$$U_i(\theta_i, z, \omega) = f(z, \theta_i) + g(z, \omega) \tag{1}$$

where $f(z, \theta_i)$ is function that equals one if $z = \theta_i$ and zero otherwise, and $g(z, \omega)$ is a function that equals one if $z = \omega$ and -1 otherwise. We might consider the first term to represent that portion of the utility that is state-of-the-world invariant. An examples of such utility might include the expressive benefit of supporting a favored candidate regardless of that candidate's chance of winning. The second term represents the benefits that are state-of-the-world dependent. For example, the utility of supporting the winning candidate, even if this is not the candidate one initially preferred.

[3] This utility function is based on the one in Austen-Smith and Feddersen (2005)

To help the committee with this decision, there are two signals about the state of the world, $c, s \in \{x, y\}$. Signal c is public information as soon as it is drawn by nature. Signal s is private information known only to player 1, who I will refer to as the "expert." Each signal has a certain probability of being correct. These are denoted p and q such that $\Pr(c = \omega) = p$ and $\Pr(s = \omega) = q$, $p, q \in (.5, 1)$. These probabilities are also public information at the start of the game.

The game proceeds as follows. In the first phase, Nature independently draws signal $c, s \in \{x, y\}$ with $\Pr(c = x) = \frac{1}{2}$ and $\Pr(s = x) = \frac{1}{2}$. It then draws ω such that $\Pr(\omega = x)$ is derived from the values of c, s, p and q. In the second stage, player one sends a message $m_1 \in \{x, y\}$ to the other players. In the third stage, each player votes $v_i \in \{x, y\}$. The state of the world is then revealed, and each player receives a payoff according to equation 1.

By assuming that voters only play weakly undominated strategies, the game becomes a relatively simple signaling game played between player one (the expert) and player two (the median voter). In these propositions I will rely on the concept of strong announcement proofness to refine the equilibrium present in the game (Matthews and Okuno-Fujiwara 1991). Briefly, a strongly announcement proof equilibrium in cheap-talk sender receiver games is a separating equilibrium that is weakly preferred by all types of senders to every other equilibrium and strictly preferred by at least one type of sender to the other equilibrium. The following propositions lay out the equilibrium:

> Proposition 1: There is always a pooling equilibrium in which the message strategy of the first player is independent of the signal and the voting strategy of the second player is independent of the message send by the first player.

Proposition 2: If the expert and the second player are of the same type, there is always a separating equilibrium in which $m = s$ and $z = x$ if $EU(x|\theta = x, c, s) \geq EU(y|\theta = x, c, s)$ and y otherwise. If $EU(x|\theta = x, c, s = y) < EU(y|\theta = x1, c, s = y)$ and $EU(x|\theta = x, c, s = x) > EU(y|\theta = x, c, s = x)$ this equilibrium is strongly announcement proof.

Proposition 3: If the expert and the second player are of opposite types, there is only a separating equilibrium in which $m = s$ if one of the following sets of conditions holds

1. $EU(x|\theta = y, c, s = y) > EU(y|\theta = y, c, s = y)$ or $EU(x|\theta = x, c, s = x) < EU(y|\theta = x, c, s = x)$.
2. $EU(x|\theta = y, c, s = x) < EU(y|\theta = y, c, s = x)$ or $EU(x|\theta = x, c, s = y) > EU(y|\theta = x, c, s = y)$.
3. $EU(x|\theta = x, c, s = y) < EU(y|\theta = x, c, s = y) \text{and} EU(x|\theta = x, c, s = x) > EU(y|\theta = x, c, s = x)$ and if $EU(x|\theta = y, c, s = y) < EU(y|\theta = y, c, s = y) \text{and} EU(x|\theta = y, c, s = x) > EU(y|\theta = y, c, s = x)$

If the third condition holds, this equilibrium is strongly announcement proof.

Proposition 1 states the familiar babbling equilibrium in cheap-talk games in which in which the expert sends an uninformative signal and the median voter ignores it. I will not prove this proposition directly, but rely on the intuition explained in section Farrell and Rubin (1996). Proof of the other propositions is left in Appendix A,[4] but

[4]This and all other appendices to this paper are available at `http://www.princeton.edu/~cdmyers/DCC_app.pdf`.

I will try to give some sense of the intuition behind them here. At the most basic level the actions of the players are determined by the strength of the signals they receive. If a signal is very strong, the expected utility of choosing the policy suggested by the signal outweighs the guaranteed utility a player gets from choosing his ex ante preferred policy. This can also happen if two signals point in the same direction and the combination of the two signals is very strong. However, if two strong signals point towards opposite policies they may cancel each other out, returning the player to his ex ante preferred policy.

Given this, there are four types of signal strength combinations that can exist.

I: If both signals are weak, participants will ignore their content since even if both point in the same direction there will still be sufficient uncertainty about the state of the world that the expected utility of choosing the indicated policy will not outweigh the utility gained by choosing the ex ante preferred policy. In this case a separating equilibrium exists, but is not Strongly Announcement Proof.

II: If the public signal is very strong relative to the private signal, both types of players will ignore their ex ante preferred policy and vote for the policy indicated by the public signal; they will also ignore the private signal (and any message about it) since it is not sufficiently strong to change their opinion about the likely state of the world. Again, a separating equilibrium exists that is not Strongly Announcement Proof.

III: If the private signal is very strong relative to the public signal both types will prefer the policy indicated by the private signal. In this case the expert will be able to honestly communicate the private signal to the median voter regardless of the median's type because both types have the same preferred policy and both types know they have the same preference.

IV: If the private and public signals are both relatively strong and are of relatively equal strength both types prefer their ex ante preferred policy if the signal disagree, but if the signals agree both prefer the policy indicated by the two signals. If the expert and median are of the same type the expert will honestly communicate his private signal to the non-expert because they share a preferred policy and they both know they share a preferred policy. If the expert and median are of different types and the public signal is opposed to the median's preference the expert can not honestly communicate his private signal, since given a private signal that is the same as the median' preference they do not share a preferred policy and both know that they do not share a preferred policy.

More formally, the propositions lay out when a separating, or informative, equilibrium exists in addition to the always present pooling, or uninformative, equilibrium, as well as when it is plausible that the separating equilibrium will be consistently selected instead of the pooling equilibrium. Proposition 2 states that when the expert and the pivotal voter share the same utility function honest revelation is always in equilibrium. This equilibrium is strongly announcement-proof when the private signal is sufficiently strong relative to the public signal that its contents can cause type x voters to follow the private signal regardless of with the public signal says, types III and types IV when the public signal points in the opposite direction of the ex-ante preferred policy. In cases I, II or some cases in IV, the message sent by the expert is irrelevant, so any messaging strategy is in equilibrium.

Things become a bit more complicated when the pivotal player is the opposite type of the expert. Parts 1 and 2 of Proposition 3 state that honest revelation is possible in case I and II, but that this equilibrium is not strongly announcement proof. Part three states that honest revelation is possible when the private information is sufficiently strong relative to the public signal that both player prefer to follow it regardless of the public information, case III. Otherwise (case IV) honest revelation is impossible, and only the babbling equilibrium exist.

Of course, these general statements are meaningless until they are expressed in terms of the utility function described in equation 1. Doing so requires deriving eight different expected utilities. For brevity, I will relegate this step to the appendix[5] and present the results graphically.

Figure 1 defines four regions in the signal-strength space. Table 1 list the equilibrium that are present in each region under each set of circumstances. Where the equilibrium is strongly announcement proof I do not list the pooling equilibrium. Region I represents case I discussed above; both signals are weak and even if they point in the same direction both types prefer to follow their type. In region II the public signal is sufficiently strong relative to the private signal that booth type prefer to enact the policy indicated by the public signal regardless of the private signal - this is case II. Region III corresponds to case III where the private signal is sufficiently strong that both types prefer to enact the signal indicated by the private signal so the separating equilibrium is strongly announcement proof.

In region IV the predictions are a bit more interesting. Given a public signal that agrees with the median voter's type she always prefers to vote for here type. However,

[5] Available at http://www.princeton.edu/~cdmyers/DCC_app.pdf.

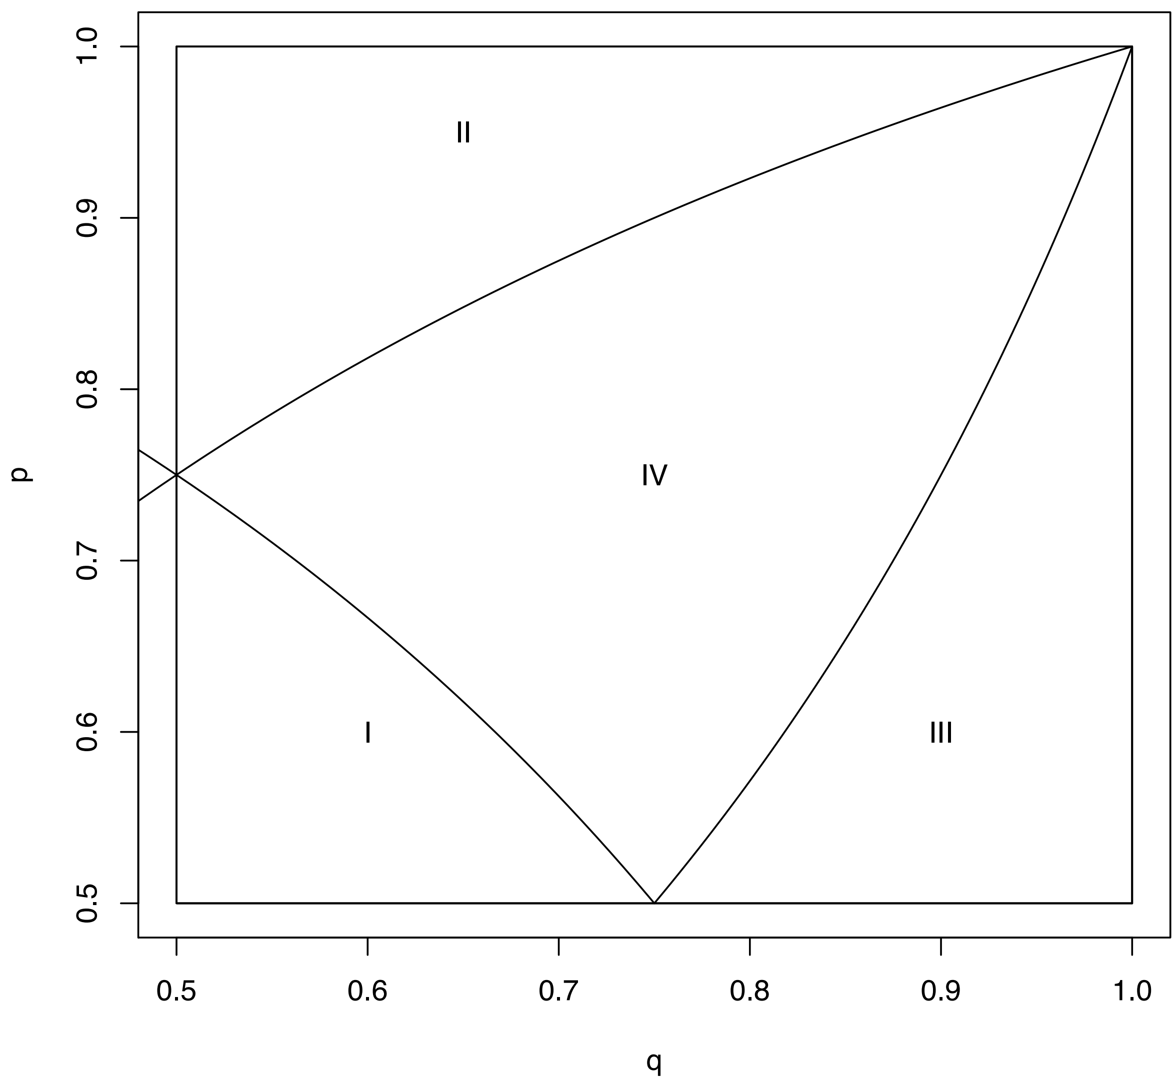

Figure 1: Signal Strength Regions

	$\theta_2 = x$		$\theta_2 = y$	
Region	$c = x$	$c = y$	$c = x$	$c = y$
I	Pooling and Separating	Pooling and Separating	Pooling and Separating	Pooling and Separating
II	Pooling and Separating	Pooling and Separating	Pooling and Separating	Pooling and Separating
III	Separating	Separating	Separating	Separating
IV	Pooling and Separating	Separating	Pooling	Pooling and Separating

Table 1: Equilibriums by Region and Signal Type

if the public signal disagrees with the median voter's type, she will vote against her signal only if the private signal agrees with the public signal. In this case, the expert will honestly reveal his signal if the median voter is of the same type, but will be unable to do so if the median voter is of a different type.

2.0.1 Hypotheses From the Formal Model

The testable hypotheses of this formal model focus on whether the predictions for region IV are accurate.

H_f1: If the signal strengths of private and public information fall within region IV, private information will only be shared if the median voter on the committee has the same preference than the expert.

H_f1: If the signal strengths of private and public information fall within region IV, the median voter does not have the same preference as the expert, and the private information has different decision implications than the public information, the committee will reach a different decision than it would if all information was public.

3 Experimental Design

This experiment tests the prediction of the model described above alongside the prediction of the common knowledge effect literature. Since the experiment is intended to test the group dynamics predictions of the psychological literature as well as the strategic predictions of the game-theoretic model, it must be a hybrid of two experimental traditions. Testing the game-theoretic predictions requires inducing preferences by making participants' earnings depend on their performance during the experiment. These preferences will be varied as in the model described above: the participant with "private" information will either share the preference of the majority of the committee, or will be the only committee member with her preference. Incentivized preference are rare in psychological experiments, and no study in the common knowledge effect literature does so. Testing the group dynamics predictions of the psychological literature requires that the "privateness" of the private information be varied. I will do this in two ways: by making the private information public, and by identifying the participant who knows the private information to the other participants.[6] Combined, these two manipulations identify the independent effects of these two processes and show their interaction.

Participants are told that they are taking part in an experiment on group decision making, and that their earnings will depend on the decision reached by the group. Participants are then placed in groups of three tasked with choosing one of two policy options. I use three different decision scenarios. For clarity I will limit discussion in this document to one of them: a political club that must choose to endorse one

[6]Testing these predictions also requires that participants be allowed to deliberate in a natural way; economics experiments testing models of information sharing rarely include actual conversation.

of two candidates in an upcoming election.[7]Each group member has an affinity for one candidate and gains some expressive benefit from supporting him. However, all members want the club to endorse the winning candidate so that they can have some influence over policy during his administration. Committee members are incentivized to hold the appropriate preferences by making their earnings from the experiment dependant on the decision the committee makes. In line with the model described above, each participant receives five dollars if their favored candidate is selected by the group and ten dollars if the winning candidate is selected by the group.

The participants are given biographies of the candidates which contain several pieces of neutral information along with one or two pieces of information that are unambiguously positive for one or the other candidate. One of these pieces of information will be "private" information; the number and identity of the deliberators who receive this information is experimentally manipulated.

3.1 Experimental Conditions

The deliberator (or deliberators) who receive the "private" information is manipulated in two ways, corresponding to the game-theoretic and psychological theories of information sharing. Consider item of information a. First, following the game theoretic models, the composition of the committee is varied so that the committee member who knows a has an incentive to share or hide the information. In the "Information Holder In Majority" condition, one other member shares the preference of

[7]The other two scenarios are a town that must decide whether to expand its sewage treatment facility to accommodate a new factory that might be built and a town that must decide between two plowing contractors, one expensive and one inexpensive, in anticipation of either light or heavy snowfall in the coming winter. Both are detailed in Appendix B available at `http://www.princeton.edu/~cdmyers/DCC_app.pdf`.

the member that knows a. In the "Information Holder Not In Majority" condition the information is given to the committee member who is in the minority. This simple manipulation, which preserves the relative balance of the committee, is sufficient to test the model described above.

Secondly, the common knowledge effect is manipulated in two ways. The first is the distribution of information; a is either known to all members of the committee before deliberation or it is known by only one member of the committee before deliberation. However, this manipulation can not be fully crossed with the majority-minority manipulation, since information that is held by all members of the committee is clearly held by a member of the majority of the committee. Thus I manipulate the common knowledge effect in a second way, by identifying the deliberator who knows a. As discussed above, identifying an experimental participant who has unique knowledge of some information essentially eliminates the common knowledge effect (Stasser, Vaughan and Stewart 2000). Even if all participants do not know a piece of information, telling them that that information exists (and is known by one of their discussion partners) is sufficient for that information behave as though it were common knowledge. This "Expert IDed" manipulation allows for the complete crossing of the common knowledge effect manipulation and the preference manipulation.

The result is five experimental conditions, outlined in Table 2. Inferences are drawn by comparing deliberations where the same item of information is "private" - only the participant (or participants) who have access to that information will be changed. Specifically, the following relationships are predicted by the theories outlined above:

- From the Game Theoretic Model: Greater information sharing in cell 2 then

		Information Holder in Majority?	
		No	Yes
Information is “Common Knowledge?”	No	1	2
	Yes - Holder is IDed	3	4
	Yes - Known by all	X	5

Table 2: Five Experimental Conditions

cell 1, cells 4 and 5 than cell 3.

- From the Psychological Model: Greater information sharing in cell 3 than cell 1, cells 4 and 5 than cell 2. Same amount of information sharing in cells 4 and 5.

- Combined theory: Greater information sharing in cells 4 and 5 than in cells 2 and 3, greater information sharing in cells 2 and 3 than in cell one. Same amount of information sharing in cells 4 and 5.

The combined theory makes no prediction about the relationship between cell 2 and cell 3. However, this relationship is still of interest, since it can suggest which theorized impediment to information sharing causes the greatest problem. If there is greater information sharing in cell 2, the common knowledge effect is not as great an impediment as the strategic incentives of the game theoretic model. If the relationship runs the other way, then the opposite is true.

3.2 Information

The signals described in the formal model are represented in the experiment by items of information about the two candidates. This experimental set-up requires that the signals have strengths that place the deliberation is region IV of Figure 1. Roughly speaking, this requires that both signals be fairly strong and of similar strength. In the experiment, these signals are pieces of information about the two candidates that favors the chances of one candidate over the other. Each signal contains positive information about one of the candidates, implying that that candidate has a greater chance of winning the election than his opponent (and will thus resemble the signals discussed in the formal model). The items of information are advanced notice about an upcoming newspaper endorsement and advanced access to preliminary fundraising totals. The information is embedded in candidate biographies amidst several pieces of neutral information that does not favor either candidate. For instance, the candidates are portrayed as having similar past experience in public office, having attended similar colleges, having similar families, and equal bases of support.

In conditions where there is an expert one of the manipulated pieces of information is disclosed only to the expert. This information is presented to the expert as private information that the other group members are unaware of. The piece that is shown only to the expert is rotated, so that the strength of that information when it is private can be compared to when it is public knowledge. Since the private information must be something that the other committee members might plausibly be unaware of, both pieces of manipulated information are somewhat non-public in nature.[8]

The simplest way to ensure that the items of information had the proper "signal strength" would be to explicitly attach a particular value to each signal. For example, the signal might say "Jon Smith will be endorsed next week by the *Springfield Shopper*, an endorsement that gives him a 80 percent chance of winning the race." However,

[8]For example, the fact that one of the candidates has a law degree while the other does not is not a fact that could not plausibly be private information. However, the fact that one is about to be endorsed by a major newspaper might be know to only a few people.

such an explicit translation of events into probabilities seems unrealistic. More importantly, natural deliberation about the candidates is unlikely after the participants are presented with such a stark statement. This would both harm the external validity of the experiment and make testing the psychological theories impossible. Instead, I establish the strength of the two signals by providing historical information about the value of the signals (i.e."The candidate receiving this endorsement has won in three of the last four elections for county commissioner.")

The background information for the candidate endorsement task is in Figure 2. The private and public signals, enclosed in brackets in the text, are interspersed in the neutral biographies. Here, the signals are opposed; the public signal favors Jones while the private signal favors Davies. In the "known to all" condition, all deliberators would see both pieces of information, with the private signal changes slightly so that it is clear that all committee members are aware of it. In the conditions where one participant has private information the private signal would be replaced with a neutral statement that indicates that the information (in this case the newspaper endorsement) is unknown, but that if it were know it would have a particular strength. In all conditions, the valance and the signal that is private is randomized.

3.3 Experimental Procedure

Each session involves nine participants, five women and four men or four women and five men. The session will consist of three rounds. In each round the participants are rematched into new groups of three, and given a new decision task as described above.[9] No subject is ever be in the same group with any other subject more than

[9] As noted above, this task will not always be candidate endorsement.

In this task you must decide to endorse one of two candidates for county commissioner. Each person wants to endorse the candidate who will go on to win the election. However, each person also leans towards one of the candidates. ***You lean towards Mark Jones.***

Candidate 1: Mark Jones

Mark Jones was born in Mercerville, a small town in the western portion of the county. His mother, a school teacher, impressed upon him from an early age the importance of education. He graduated from Mercerville High School and attended American University in Washington DC on an Air Force ROTC scholarship. He served five years in the Air Force, including combat missions during the first Gulf War.

After leaving the Air Force he attended Georgetown Law School and returned to the area to set up practice in Springfield, the major city in the county. Jones has served two terms on the County Commission. He ran for state legislature last year, losing in the primary to former mayor of Springfield. Jones has been married for 14 years to fellow lawyer Diane Reading, with whom he has three children.

Jones launched his campaign for county commissioner last February, and has already shown strong fundraising skills. He has used his connections on the county commission to raise money. Jones is supported by a number of local politicians, including the town's Prosecutor and the former County Commissioner. Jones's campaign slogan is "Preserving a Better Tomorrow." He plans on focusing on his experience on the county commission, as well as his long-standing ties to the area.

Candidate 2: Tom Davies

Tom Davies was born in Cedar Springs, a blue-collar suburb of Springfield. His father, a machinist, taught him the value of hard work and dedication. After graduating from Springfield High School he enlisted in the Army, serving for 8 years and rising to the rank of Staff Sergeant. During his service he led troops in Bosnia and Haiti.

Davies earned a Bachelor's degree while in the Army from George Washington University. After leaving the service he earned a law degree from the University of Michigan Law School and then moved back to Springfield and set up his own practice. He won reelection to the City Council twice, resigning last year to mount an unsuccessful campaign for Mayor. Davies is married to Amanda Dyson, a teacher at Springfield High School, and has two children aged 8 and 10.

Davies announced his run for county commissioner in March, and had been quietly raising a considerable number of donation pledges for months before then. {In fact, the committee has learned that Davies has raised $50,000 to his opponent Mark Jones's $25,000. Four of the last six candidates who enjoyed this sized fundraising advantage in recent county commission elections have gone on to win.} {You have recently learned that the Springfield Shopper, the highest circulation daily newspaper that reaches about three quarters of the county's population, has decided to endorse Davies. The candidate receiving this endorsement has won in two of the last three elections for county commissioner. *The other people are not aware of this information*} Davies is running on a platform of "Our county, Our tomorrow." His campaign focuses on his knowledge of Springfield's people and problems, as well as his military service.

Figure 2: Sample Candidate Biographies

once. Each session provides nine group-level observations per session

Participants begin the experiment by reading the information provided to them about the decision task; in the case of the candidate selection task, this is the candidate biographies. After reading the information, I administer a pre-deliberation questionnaire will make sure that participants have learned the information they are supposed to use in deliberation and measure their beliefs about the chances of each candidate before discussion begins. Next, participants deliberate and come to a decision by majority rule. After deliberation, they are asked to complete a post-deliberation survey. After all sessions are finished participants complete a demographic questionnaire and will then be paid. Deliberation is voice recorded.

3.4 Results

I begin with an analysis of the decisions made by groups. Ideally, the information environment in which a group deliberates should not effect the decision made by the group. In other words, groups that deliberate when all information is public should reach the same conclusion as when some information is private and that information is held by a group member who is in the minority. However, the theory discussed above suggests that this will not happen.

Table 3 shows the percentage of groups in each condition that reach the "correct" decision. Note that in this context "correct" means simply the decision that maximizes the earnings of the members of the majority - in other words, the decision we would expect the committee to make given complete information. As expected, all groups in conditions 4 and 5, the conditions that mimic complete information, make the "correct" decision. In condition 2, where strategic incentives suggest that infor-

		Information Holder in Majority?	
		No	Yes
Information	No	60 (n = 5)	80 (n = 5)
is "Common	Yes - Holder is IDed	80 (n = 5)	100 (n = 7)
Knowledge?"	Yes - Known by all	X	100 (n = 5)

Table 3: Percentage of "Correct" Decisions

mation should be shared but psychological theories suggest incomplete sharing, only four out of five groups reach this decision. The same number come to the "correct" decision in condition 3, where psychological theories predict full sharing but a strategic incentive makes it hard for the participant with secret information to share it. Finally, in the condition where both psychology and game-theory predict incomplete sharing only three out of five groups reach the correct decision. These differences can be attributed to the shifting information environment, though the small sample sizes mean that these differences are not statistically significant. These results suggest that altering the distribution of information can change the decisions groups make even if the information itself does not change. Further, they suggest that both theories as to why this might happen are correct.

Can we see these effects in the content of group discussion? To test this, I look at the number of times each item of information was mentioned in group discussion. Table 4 shows the bias in favor of discussing public vs. private information in each of the experimental conditions. The statistic shown is the average across all observations in each condition of the following:

$$\frac{\text{Number of Mentions of Public Information} - \text{Number of Mentions of Private Information}}{\text{Total number of Mentions of Both Items of Information}}$$

		Information Holder in Majority?	
		No	Yes
Information is "Common Knowledge?"	No	.33	-.22
	Yes - Holder is IDed	-.26	0
	Yes - Known by all	X	-.13

Table 4: Weighted Difference in Mentions Between Public and Private Information

Roughly, this shows the bias in favor of discussing public information weighed for the total amount of discussion in each group. A score of 1 indicates that the private information was not mentioned, with a 0 indicates that both items were mentioned an equal number of times.

Here, the results are mixed. Consistent with theory, groups in condition 1 show a strong tendency to discuss public information more than private information while groups in condition 4 showed no bias in their discussion. However, groups in conditions 2 and 3 tended to discuss private information more than public information, contrary to theoretical predictions.

Given the small sample sizes, these results are highly preliminary. Additional analysis will include more careful coding of the discussion to provide a more precise measure of how often each item of information was discussed, as well as the manner in which is it was discussed.[10] I will also make use of measures of altruism, need for cognition, and demographics to see the impact of these factors on the willingness to share information.

[10]For example, a group might discuss whether information from a member with a strategic incentive to lie could be trusted. This would show up as a large number of mentions of that item of information without the information actually being conveyed in deliberation.

4 Conclusion

Like all normative theories of government, the theory of deliberative democracy aims to affect how collective decisions are made. More importantly, theories like deliberative democracy grant legitimacy to collective decisions that meet the procedural and substantive demands of the theory. But unless we understand the how these theories function in the real world, we risk granting legitimacy to decisions that do not deserve it. This experiment tests whether a key part of deliberation, the exchange of information, functions sufficiently well to legitimize decisions made by deliberative processes.

This question is a key point of contention between positive and normative theorists of deliberation. Positive theorists argue that the strategic incentives embedded in communication make it difficult for deliberation to function properly (Landa and Meirowitz 2003). Normative theorists, when they reply to these claims, tend to argue that the sparse, stylized models of game-theorists can tell us little about how actual people deliberate. The solution to this conflict is to bring empirical evidence to bear. By embedding strategic incentives in a more natural communication environment, I hope to provide such data.

Initial results suggest that deliberation can not be counted on to make efficient use of information is all cases. Moreover, this inefficiency is biased in that information held by certain parties is less likely to me used than other information. In particular, groups make different decisions when important information is held by fewer members, and this effect increases when the people who hold the information is a preference minority.

References

Austen-Smith, David. 1990. "Information Transmission in Debate." *American Journal of Political Science* 34(1):124–152.

Austen-Smith, David & Timothy J. Feddersen. 2005. Deliberation and Voting Rules. In *Social Choice and Strategic Decisions: Essays in Honor of Jeffrey S. Banks*, ed. David Austen-Smith & James Dugan. New York: Springer.

Benhabib, Seyla. 1996. Toward a Deliberative Model of Democratic Legitimacy. In *Democracy and Difference: Contesting the Boundries of the Political*, ed. Selya Benhabib. Princeton: Princeton University Press pp. 67–94.

Cai, Hongban & Joseph Tao-Yi Wang. 2006. "Overcommunicating in Strategic Information Transmission Games." *Games and Economic Behavior* 56(1):7–36.

Cohen, Joshua. 1989. Deliberation and Democratic Legitimacy. In *The Good Polity*, ed. Alan Hamlin & Phillip Petit. Boston: Blackwell pp. 17–34.

Crawford, Vincent & Joel Sobel. 1982. "Strategic Information Transmission." *Econometrica* 50:1431–1451.

Cruz, Michael G., Franklin J. Boster & Jòse I. Rodrìguez. 1997. "The Impact of Group Size and Proportion of Shared Information on the Exchange and Integration of Information in Groups." *Communication Research* 24(3):291–313.

Dickhaut, John W., Kevin A. McCabe & Arijit Mukherji. 1995. "An experimental study of strategic information transmission." *Economic Theory* 6(3):389–403.

Farrell, Joseph & Matthew Rubin. 1996. "Cheap Talk." *Journal of Economic Perspectives* 10(3):103–118.

Gigone, Daniel & Reid Hastie. 1993. "The Common Knowledge Effect: Information Sharing and Group Judgement." *Journal of Personality and Social Psychology* 65(5):959–974.

Gutmann, Amy & Dennis Thompson. 1996. *Democracy and Disagreement*. Cambridge: Cambridge University Press.

Gutmann, Amy & Dennis Thompson. 2004. *Why Deliberative Democracy?* Princeton: Princeton University Press.

Krishna, Vijay & John Morgan. 2001. "A Model of Expertise." *The Quarterly Journal of Economics* 116(2):747–775.

Landa, Dimitri & Adam Meirowitz. 2003. "Game Theory, Information, and Deliberative Democracy." Princeton Typescript.

Matthews, Steven A. & Masashiro Okuno-Fujiwara. 1991. "Refining Cheap Talk Equilibrium." *Journal of Economic Perspectives* 55(2):247–273.

Meirowitz, Adam. 2006. "Designing Institutions to Aggregate Preferences and Information." *Quarterly Journal of Political Science* 1:373–392.

Parks, Craig D. & Cowlin Rebecca. 1995. "Group Discussion as Affected by Number of Alternatives and by a Time Limit." *Organizational Behavior and Human Decision Processes* 62(3):267–275.

Stasser, Garold & Dennis Stewert. 1992. “Discovery of hidden profiles by decision-making groups: Solving a problem versus making a judgment.” *Journal of Personality and Social Psychology* 63(3):426–434.

Stasser, Garold, Sandra I. Vaughan & Dennis D. Stewart. 2000. “Pooling Unshared Information: The Benefits of Knowing How Access to Information Is Distributed among Group Members.” *Organizational Behavior and Human Decision Processes* 82(1):102–116.

Stasser, Garold & William Titus. 1985. “Pooling of Unshared Information in Group Decision Making: Biased Information Sampling During Discussion.” *Journal of Personality and Social Psychology* 48(6):1467–1478.