

Interests, Information and Minority Influence in
Deliberation

C. Daniel Myers

Department of Political Science, University of Minnesota

Minneapolis, MN 55455

cdmyers@umn.edu.

May 16, 2016

Abstract

The ability of citizens to advance arguments and have them considered fairly is essential to the equality and epistemic quality of deliberative institutions. This paper builds on game-theoretic models of strategic information transmission to offer a theory of how the interests that deliberators have in the outcome of deliberation can reduce the influence of some citizens in deliberative process. Specifically, this paper argues that the influence of an argument depends on whether the deliberator making the argument is in the majority or minority in terms of her interests in the outcome of deliberation. An arguments made by a member of the minority will be less influential than the same argument made by a member of the majority. We offer the first empirical test of this kind of model in realistic deliberative conditions using a laboratory experiment and a field experiment, and find support for this theory.

Keywords: Deliberation, Information Transmission, Interests, Information, Democracy

Supplementary materials including proofs of propositions, experimental instructions, and supplementary analyses are available in appendices to the online edition. All data and code necessary to replicate the analyses in this article are archived in the Journal of Politics Dataverse <https://dataverse.harvard.edu/dataverse/jop>. The studies described in this article were reviewed and approved by the Institutional Review Board at Princeton University. The research received financial support from the Mamdouha S. Bobst Center for Peace and Justice at Princeton University and the Robert Wood Johnson Scholars in Health Policy program.

One of the core elements of deliberation is the exchange of information. Whereas aggregative versions of democracy hold that legitimacy lies in the fair counting of votes for and against a decision, deliberative democrats hold that a decision is only legitimate if it is informed by a fair discussion (Cohen 1989, Gutmann and Thompson 1996). “Citizens and their representatives are expected to justify the laws they would impose on one another by giving reasons for their political claims and responding to others reasons in return” (Thompson 2008, pg. 498). In addition to the normative value of discussion, information-sharing is helpful for many of the practical benefits that proponents hope deliberative institutions will provide, such as higher quality policy (Estlund 1997, pg. 173) or a greater appreciation of the views of the opposing side.

Given the centrality of information to deliberation’s value, it is not surprising that many critiques of small group deliberation are rooted in concerns about its ability to foster the free and fair exchange of information. Consider, for example, critics who are concerned that deliberation will sideline the voices of women and minorities. These critiques argue that institutional features of the deliberative environment (e.g. Karpowitz and Mendelberg 2014), persistent racism, sexism and other biases (e.g. Sanders 1997), and the tendency of deliberation to privilege certain kinds of speech (e.g. Young 1996), make information known by members of marginalized groups less influential in deliberation than information known by members of less marginalized groups. This is normatively troubling because it means that individuals are treated unequally in deliberation, and also because the quality of group decision making is reduced because information that is relevant to the group’s decision is excluded or not given proper weight.

However, gender and racial prejudice are not the only biases that can harm the quality of deliberation. In this paper we focus on another bias in group deliberation - bias against arguments made by deliberators who are in the minority in terms of their interests in the outcome of deliberation. While research on racial and gender bias in deliberation is motivated by the literature on race and gender in politics, the current paper is motivated by game-theoretic models of information sharing in deliberation, as well as renewed attention to the

role of interests in deliberation on the part of normative theorists (Mansbridge et al. 2010). These game-theoretic models predict that when deliberators have conflicting interests in the outcome of deliberation, deliberating groups will discount information contributed by those whose interest is in the minority (henceforth the “interest minority”). They do this because deliberators in the interest minority, unlike deliberators in the interest majority, have an incentive to misrepresent information in order to lead the group to a decision that is in line with the minority’s interests. However, the result of this bias is that all information contributed by members of the interest minority is discounted, even when this information is an important and honest contribution to deliberation.

We test for the presence of this bias using two experiments. The first is a laboratory experiment that allows for precise control of the information deliberators have about the group decision as well as the interest each has in the outcome of that decision. The second is a field experiment that gives up some experimental control to test this theory in a more externally valid manner. In both experiments we test for bias against arguments made by members of the interest minority using a method that is common in the “hidden profile” literature in social psychology (Stasser and Titus 2003), but is new to political science. We adapt this technique by randomly assigning deliberating groups to two conditions, one where an item of information is known only by a member of the interest majority and one where the same item of information is known only by a member of the interest minority. We then compare the influence of this item of information across these two conditions to test whether the interests of the person who knows an item of information affect how influential that item is in discussion. In both experiments we find that an item of information has more influence in deliberation when it is known by a member of the interest majority than when it is known by a member of the interest minority.

The broader contribution of this paper is to demonstrate that interests matter for the conduct of small group deliberation. The idea that deliberators’ interests influence the information that they choose to communicate and the way that they interpret the information communicated by others is not new (Landa and Meirowitz 2009). Indeed, it is the basis

for a large game-theoretic literature on information in deliberation. However, these insights remains untested in realistic deliberative conditions; as a result, this literature remains largely ignored by most normative and empirical scholars of deliberation.¹ The empirical tests in this paper offer evidence that interests affect the way that deliberating groups use information to make decisions, and will hopefully spark greater interest in the insights of the game-theoretic literature among scholars of deliberation.

On a practical level, these results reinforce the importance of providing information that is perceived as coming from credible neutral sources, as is done in some mini-publics (e.g. Deliberative Polls). However, the provision of information from neutral, credible sources can be costly, and may be difficult from some kinds of mini-publics to achieve. Further, the provision of neutral information may be impossible for some important types of information, such as contested areas of knowledge where no experts who are perceived as credible and neutral, or the case of situated knowledge that is known only by deliberators themselves (Young 1996).

1 Information and Deliberation

The sharing of information plays a key role in most deliberative theories of democracy (Goodin 2012, Ch. 5). Here we take a broad definition of information, encompassing arguments, facts, stories of personal experience, moral claims, or other forms of dialogue that support a particular policy. Deliberative democrats argue that policy is only legitimate if it is preceded by a debate where information supporting various policy options is exchanged, and in particular where proponents of particular policies offer information supporting those policies that other members of society might be reasonably expected to accept (Cohen 1989, Gutmann and Thompson 1996). Further, deliberators must have an equal chance of having their arguments fairly considered by others (Knight and Johnson 1997). The desired outcome of properly structured deliberation is therefore that “no force except that of the better argument is exercised” (Habermas 1975, pg. 108). This focus on argu-

¹Though see see Myers and Mendleberg (2013).

ments’ “force” alone ensures that all citizens have an opportunity to influence deliberation (Gutmann and Thompson 1996, Ch. 3).

Whether the influence of arguments in deliberation is determined by their force alone is important to judging two key qualities of deliberative institutions: deliberation’s epistemic value, or whether they make good decisions (Estlund 1997), and deliberation’s equality, or whether all can present arguments and have them fairly considered (Knight and Johnson 1997). Deliberation has epistemic value because it requires people to consider information that they would not have otherwise considered, provided that they “evaluate arguments fairly, irrespective of the identity of the person or the size of the group offering them” (Estlund 1997). While the epistemic value of deliberation promises better policy, equality is essential to deliberation’s democratic legitimacy. Deliberation that gives less-than-full consideration to arguments offered by some groups or individuals violates this rule of equality, and thus provide less democratic legitimacy. Thus, a key task of empirical scholars of deliberation is to identify biases that jeopardize these qualities in actual deliberative institutions.

1.1 Strategic Information Transmission and Deliberation

One such bias that has attracted little attention from normative and empirical scholars of deliberation is the way that conflicting interests can produce biased consideration of arguments advanced by members of the interest minority. We define “interests” as preferences for the outcome of deliberation that are unlikely to be revised through the deliberative process itself. While this includes material self-interest, it also includes other-regarding preferences and other sources of what Gutmann and Thompson (1996, pg. 25) refer to as “persistent moral disagreement,” such as religious doctrines or ideologies.² Deliberators can and generally do have multiple interests, some of which may be held in common with others and others of which are not shared. We define an “interest minority” as an individual or group with a

²This is similar to Mansbridge et al.’s (2010) definition of interests as denoting those preferences whose transformation “requires transforming oneself, including transforming one’s identity,” which they note is unlikely to be caused by deliberation (Mansbridge et al. 2010, fn 40), as opposed to “surface preferences, opinions, or pre-reflective understandings” that may change as one learns through the process of deliberation (Mansbridge et al. 2010, fn 15).

substantial interest that is not shared by a majority of deliberators; an interest majority is a group with a substantial interest that, while not shared by all other deliberators, is shared by a majority of other deliberators.

This bias is based on the findings of the game-theoretic literature on strategic information transmission, which argues that the influence of arguments in discussion depends on the interests of the person offering the argument, as well as how those interests compare to the interests of other deliberators. This literature's basic finding is that when all deliberators have the same interest in the decision being made, all arguments will be fairly evaluated. However, if deliberators hold conflicting interests, some deliberators may have an incentive to lie or withhold information. Given this incentive, other deliberators will view the statements of a deliberator who has an incentive to mislead with suspicion, *even when these statements are true*. This bias against interest minorities represents an important challenge for the design of deliberative institutions, particularly those prescribed for situations where interests and values are in conflict.

Most of the game-theoretic literature on deliberation is based on cheap-talk models of communication.³ Beginning with Crawford and Sobel (1982), cheap-talk models conceptualize communication as an interaction between a sender and a receiver. The sender knows some information relevant to a decision that the receiver has to make (henceforth the sender's "private information"). The sender sends a message to the receiver and the receiver then makes a decision. Importantly, both sender and receiver have an interest in the decision that the receiver makes - these interests lead them prefer some decisions over others, and the private information known by the sender would help both determine which of the decisions is in line with their interests.

Crawford and Sobel show that when both players have the same interests regarding the decision the receiver will make, then the sender can communicate private information perfectly. However, if interests are not identical, then honest revelation may not be possible. Since the sender cares about the decision that the other player will make, the sender cares

³For an exception see Dickson, Hafer, and Landa (2008)

about how the information she communicates affects that decision. Thus when a sender believes that honestly communicating her private information would lead the receiver to make a decision she does not like, she has an incentive to misrepresent her private information. Given this incentive, the receiver is suspicious of all of the sender's statements. Since it is impossible to tell the difference between a statement that honestly reports private information supporting the sender's position and a statement that dishonestly reports private information supporting the sender's position, the receiver will doubt the sender when she tells them the truth as well as when she lies. The result is that the receiver discounts information communicated by a sender who do not share his or her interests, and that such information is given less weight (often, no weight) in the receiver's ultimate decision.

Models exploring the institutional design of deliberation are interested in determining when a person who knows private information can credibly communicate that information to other deliberators; in game-theoretic terms, determining when a fully revealing equilibrium exists. As in Crawford and Sobel's two-person case, these models have shown that fully revealing equilibrium do not always exist unless all deliberators share the same interest (Austen-Smith and Feddersen 2006, Gerardi and Yariv 2007). Further, when all deliberators do not share the same interest, it is group members who are in the minority in terms of their interests who have a hard time communicating private information. Meirowitz (2007) shows that deliberators who perceive themselves to be in the minority "have an incentive to misrepresent their information," and thus all deliberators should expect misrepresentation and will rationally ignore their arguments (Meirowitz 2007, pg. 373). As a result, even when members of the interest minority communicate honestly, their arguments have little influence in deliberation.⁴

While the theoretical finding that interest minorities have less influence in deliberation has not been tested empirically, several experimental studies find support for the basic result of the

⁴Importantly, this distinguishes strategic information transmission from motivated reasoning. Motivated reasoning leads listeners to give more weight to information that reinforces their existing opinions, regardless of its source. Strategic information transmission leads listeners to give more weight to information from members of the interest majority and less weight from members of the interest minority, regardless of whether it reinforces their existing opinion.

strategic information transmission literature: that interests shape how people communicate and interpret the communication of others (Dickhaut, McCabe, and Mukherji 1995, Cai and Wang 2006). However, these studies generally involve highly stylized decisions (for example, guessing randomly generated numbers), and none involve face-to-face communication. A partial exception is Lupia and McCubbins (1998), who find that voters in a survey experiment follow the predictions of a cheap-talk model by taking into account the interests of the speaker and the audience when deciding whether to believe statements by media sources. While none of these experiments test models of strategic information transmission in the context of face-to-face deliberation, they offer general support for this idea that interests play a powerful role in shaping how people communicate, and how listeners interpret the messages that they receive.

1.2 Information in the Empirical Literature on Deliberation

The existing empirical literature on deliberation does not examine the predictions of the game-theoretic literature described above, or the specific question of whether members of interest minorities have less influence in small group discussion. Nevertheless, the findings of several studies can be reinterpreted in light of these models. Esterling (2011) finds that the degree of polarization, or shared interests, on congressional committees has a significant impact on the kind of speech used by testifying witnesses. Grogan and Gusmano's (1997) case study of deliberation about a plan to switch Medicaid from a fee-for-service system to an HMO model found that conflicting interests held by non-profits representing the patient population made their critiques of the HMO model less credible. Finally, consider two studies of deliberation about racially charged issues. Mendelberg and Oleske (2000) find far greater racial bias in a highly contentious discussion of school district boundaries than Walsh (2007) finds in interracial dialogue groups in which interest conflict is far less salient. An interests-based explanation for these findings would suggest that members of racial minority groups frequently have private information about decisions related to racial justice (e.g. situated knowledge about the experience of racial discrimination in America), but are also frequently

in an interest minority, limiting their influence when interest conflict is salient.

While the existing literature on deliberation largely ignores the effect of interest minority status on influence in deliberation, a large psychological literature exists which shows that deliberators who are in the minority in terms of their preferences or opinions (henceforth “opinion minorities”) are less influential than deliberators who are in the majority in terms of their attitudes or opinions (see Martin et al. (2008) for a review). Two mechanisms drive this. First, “opinion majority” status acts as a correctness heuristic suggesting that opinion-majority members’ views are correct simply because they are held by more people (Mackie 1987). Second, because being a member of the majority is seen as desirable in itself, opinion-minorities are pressured to conform to the views of the rest of the group, though this influence primarily affects publically expressed views, such as statements in group discussion, rather than privately expressed views, such as answers to survey questions (Moscovici 1980, Martin et al. 2008).

However, these opinion minorities are different from the interest minorities that this paper studies. In general, opinion minorities share a common interest with other members of the group, but hold a different opinion about how to best serve that interest; interest minorities have a different interest, which may lead them to also have a different opinion about the best course of action. This difference means that neither of the mechanisms that make opinion majorities more influential (or influential through different channels) than opinion minorities affect the relative influence of interest majorities and minorities. The fact that the different views held by interest minorities are based on concrete interests means that majority status is less likely to serve as a correctness heuristic; the majority holds their views because those views are in their interest, not because they are correct in some universal sense. Similarly, the fact that majority and minority status is created by interest conflict should reduce conformity pressures. Conversely, the fact that opinion minorities share a common interest with opinion majorities means that the strategic information transmission problem does not come into play - as described in the next section, full information transmission is possible when all deliberators share a common interest. Thus while the theory tested in this

paper uses the terms “majority” and “minority,” it does so in a very different sense than the minority influence literature. The model of strategic communication tested in this paper is a separate process by which some minorities are marginalized in discussion.

2 A Model of Information Use

In this section we present a model that demonstrates the difficulty that members of an interest minority have in influencing others during deliberation. The model does so by analyzing a simple situation - a deliberating group that must make a decision with a single member who knows information that is relevant to that decision. Henceforth, we will refer to this deliberator as the “expert.” This model will generate predictions about the effect of being in the interest minority on influence by comparing equilibrium when the expert is in the interest majority to equilibrium when the expert is in the interest minority. This section will discuss the model and its implications in narrative form; proof of the model’s propositions are in online Appendix A.

Consider a group deciding between two policy options by majority vote. Each member of the group has a private interest in seeing the group choose one of the two options as well as a common interest in choosing the option that matches an unknown state of the world. The private interest is strong, but not absolute; given sufficient information about the state of the world their common interest will override their private interest and they will prefer the policy that they believe matches the state of the world, regardless of whether it also matches their private interest. Assume that all members are aware of the interests of all other members, and thus of which interest is in the majority and which is in the minority. Some information about the unknown state of the world is known by all members of the group, but one deliberator, the expert, knows one extra item of information. The group engages in discussion before the vote, and during that discussion the expert sends a message about the information that she knows. After discussion the group decides on a policy option by majority vote.

To formalize this situation, consider a committee of n where n is odd that must decide

between two policy options, x and y . There is an unknown state of the world ω that matches one of the two policy options. The final decision, z , is made by majority rule voting. Each player has a preferred policy choice that defines their type (or interest) $\theta \in \{x, y\}$. These types are public information, meaning that all players know each others' type and thus all know which type is in the majority and which is in the minority. Each player derives some utility from selecting their preferred policy and some utility from selecting the policy that matches the unknown state.⁵ Specifically, each players' utility function is:

$$U_i(\theta_i, z, \omega) = f(z, \theta_i) + g(z, \omega) \quad (1)$$

where $f(z, \theta_i)$ is function that equals one if the decision matches the deliberator's type (i.e. $z = \theta_i$) and zero otherwise, and $g(z, \omega)$ is a function that equals one if the decision matches the state-of-the-world (i.e. $z = \omega$) and -1 otherwise, so that $g(z, \omega)$ carries twice the weight of $f(z, \theta_i)$. The first term, $f(z, \theta_i)$, represents that portion of the utility that is state-of-the-world invariant, based on private interests. The second term, $g(z, \omega)$, represents the utility that depends on the state of the world, based on common interests. The balance between the second and first term represents the degree to which the players share common interests; the assumption that ω is weighed twice that of θ_i creates a situation where common interests carry more weight than uncommon interests. Thus, this utility function captures an important aspect of persuasion in deliberation: some interests are held in common, while other interests are in conflict, conflict which can not be resolved through deliberation.⁶

While the state of the world ω is unknown, the committee members know some information about it. First, there is a common prior probability π that the state of the world is x , such that $\Pr(\omega = x) = \pi$. This represents the sum of all publicly available information about ω . In addition, one player, the expert, receives a signal $s \in \{x, y\}$ that suggests that

⁵This utility function is based on Austen-Smith and Feddersen (2005)

⁶See Mansbridge et al. (2010) for a discussion of the role of such self-interest in deliberation, as well as mechanisms through which this conflict might be resolved. These authors note that ideal deliberation does not necessarily require even the potential for the common good to subsume individual interests; instead, they argue that "our ideal polity is diverse and plural ... recognizing that diversity produces irresolvable conflicts in opinion and interests."

either $\omega = x$ or $\omega = y$. This signal has probability p of being correct and is informative, such that $\Pr(s = \omega) = p > .5$. While the value of s is only known by the expert, the value of p is public information. We assume that the common prior and the private signal are sufficiently predictive that so that if both point towards the same policy as matching ω , players of both types prefer that policy, while neither is sufficiently strong such that neither the common prior nor the private signal by themselves will convince players of both types to prefer the policy favored by that information by itself.⁷ The player with private information sends a message, $m \in \{x, y\}$. We can intuitively interpret these messages as saying “the private information is x ” or saying “the private information is y .” After the message is sent, the committee chooses a policy by majority vote.

By assuming that voters only play weakly undominated strategies, the game becomes a signaling game between the expert and a player of the majority type. Two kinds of equilibrium are possible in this game. This first is a separating, or fulling revealing equilibrium where the expert sends a message that matches her private information, so that:

$$m = \begin{cases} x & \text{if } s = x \\ y & \text{if } s = y \end{cases} \quad (2)$$

and the decisive voter chooses the action that maximizes her utility given $m = s$. The second is a pooling, or uninformative equilibrium, where the expert sends the same message regardless of what her private information says and the decisive voter ignores this message. For example, if the expert’s type is $\theta = x$, she might adopt the following message strategy in an attempt to deceive the receiver:

$$m_{\theta=x} = \begin{cases} x & \text{if } s = x \\ x & \text{if } s = y \end{cases} \quad (3)$$

Since m is the same regardless of the value of s , the decisive voter ignores m in equilibrium.

⁷This avoids trivial situations where the preferences of all of the players do not depend on either the prior or content the signal.

The results, using a perfect Bayesian equilibrium concept, show that a fully revealing equilibrium exists if and only if the expert is in the interest majority. This is because the expert and decisive voter are the same type, and thus prefer the same policy for all values of s , leading the expert to adopt a message strategy like that in Equation 2. However, when the expert is in the minority, the expert and decisive voter are of different types and thus prefer different outcomes for some values of s . This leads the expert to adopt a message strategy like that in Equation 3; the decisive voter responds by ignoring the message. Thus, the only equilibrium that exists is the uninformative equilibrium.

These results show that, despite being an interested party, an expert in the interest majority is always credible to all other deliberators, even members of the interest minority. This is because a controlling number of voters share the expert's interests and will thus have the same preferred outcome as the expert, conditional on knowing the same information as the expert. However, an expert who is a member of the interest minority is ignored by all other deliberators, including those who share her interests (i.e. those who are also in the interest minority). Thus, the model's prediction is not that individuals will discount information from experts whose interest differs from their own. It is rather that all individuals of all types will place greater weight on information from an expert whose interest is in the majority.

Based on this model, the experiments in this paper will test the following hypotheses about the influence of information that is known by interest minorities:

H1: Influence: Information known by members of the interest minority will be less influential than the same information when it is know by a member of the interest majority.

H2: Decision Quality: Group decisions will be less likely to incorporate information that is known by a member of the interest minority than the same information that is known by a member of the interest majority.

Importantly, these hypotheses test claims about the implications of the theoretical model for post-deliberation beliefs, but we do not advance claims about the specific conversational mechanisms that leads to these outcomes. While the model described here requires the expert

to send a message of either “ x ” or “ y ,” the logic of the pooling equilibrium is consistent with other kinds of behavior, such as remaining silent in all cases or saying “I know you’re not going to believe me anyways so it isn’t worth having this conversation” (Farrell and Rubin 1996). Such behavior still forms a pooling equilibrium so long as the behavior is the same regardless of value of the private signal. Relaxing the model’s restrictive signal space would allow for pooling equilibrium consistent with almost any speech behavior, including silence, making hypotheses about the exact content of the pooling equilibrium beyond the scope of this model.

3 Study 1

To test these hypotheses, this paper reports two experiments which feature face-to-face discussion in small groups. Deliberators in these groups have conflicting interests and private information about the decision that the group must make. The first study is a laboratory experiment which randomly assigns interests to subjects using cash incentives and randomly assigns one subject to know additional information about the group’s decision. These kinds of control may give rise to external validity concerns; the second study, a field experiment, is designed to address these concerns.⁸

3.1 Experimental Design

Subjects in the experiment were placed in groups of three. Each group had to make a decision with two options. One of these options would turn out to be better conditional on the unknown state of the world, but the better choice was not known by subjects until after the group made its decision. Payoffs, listed in Table 1, were structured to give subjects two interests: a common interest in selecting the option that would turn out to be better conditional on the unknown state of the world, and a private interest that did not depend

⁸This experiment contained an additional manipulation, the results of which are reported in SELF CITE, in which some subjects who knew uncommonly-known information were identified prior to deliberation. There was no interaction between this manipulation and the manipulation reported here, so we ignore it for the purposes of this paper.

on the unknown state of the world.⁹ All subjects earned \$10 if their group chose the option that turned out to “match” the unknown state of the world. However, each subject was told that they “leaned” towards one of the two options, and would earn \$5 if the group chose the option that they leaned towards, regardless of whether that option ended up matching the state of the world. Thus, if subjects knew for certain which option matched the state of the world, they would choose that option. However, with sufficient uncertainty, they would favor the option that would give them the guaranteed \$5.¹⁰ These incentives gave deliberators distinct interests while maintaining an orientation towards seeking a common solution to a problem. Two members of each group had the same private interest (i.e. leaned towards the same option) and were thus members of the group’s interest majority. The third member of the group leaned towards the other option and was the group’s interest minority.

[Table 1 about here.]

After being placed in groups, subjects read briefing materials about their group’s decision task. This background information was a narrative that contained several items of neutral information and one or two pieces of non-neutral information. Each of these non-neutral items clearly indicated that one of the two options was more likely to match the state of the world; in half of the groups both items indicated that the same option was more likely to match the state of the world, while in half of the groups each item indicated that a different option was more likely to match. One of these items was in the briefing materials of all group members (henceforth the *public item*), thus standing in for the common prior probability in the model. The second item was in the briefing materials of only one group member

⁹The use of cash incentives have a number of advantages, most importantly that they allow us to directly and precisely manipulate subjects’ interests. This advantage comes with a number of shortcomings. Perhaps most importantly, cash incentives manipulate material self-interest, as opposed to other forms of interest that might conflict in deliberation. We attempt to address these concerns in Study 2, where we manipulate the distribution in groups of subjects’ pre-existing interests instead of manipulating interests directly.

¹⁰Specifically, a risk-neutral actor will prefer the option that they lean against if the probability that this option is the correct option is greater than .75. Let x be the option that the actor leans towards, y the other option, and p be the probability that y is correct. Then $EU(x) = 5(p) + 15(1 - p)$, $EU(y) = 10(p) + 0(1 - p)$, and thus $EU(x) = EU(y)$ if $p = .75$. As is described below, we provide explicit instructions that the two items of information were independent and that each item was correct $\frac{2}{3}$ of the time. This means that if both items favored the same option that option would have a .8 probability of being correct, but neither item by itself is sufficiently predictive to cause a subject to vote against their leaning.

(henceforth the *private item*), representing the private signal in the model. Each item was accompanied by information which indicated that the item correctly predicted the better option $\frac{2}{3}$ of the time.¹¹ Each item also clearly indicated who knew that information - items that appeared in the background information of all members of the group indicated that “the committee has learned that,” while those known by only one member of the group included a phrase like “you have learned” and “others are unaware of this information.” Subjects completed a pre-discussion comprehension quiz ensuring that they properly understood this information; those who gave incorrect answers were corrected in private by the experimenter before discussion began.

The experiment consisted of three conditions, defined by which deliberator knew the private item before discussion began. In the Expert-in-Majority condition the private item appeared in the background information of one of the members of the interest majority. In the Expert-in-Minority condition the private item appeared in the background information of the member of the interest minority. In the Known-by-All condition the private item of information appeared in the background information of all members of the group, and was thus indistinguishable from the public item of information.¹² This allows the Known-by-All condition to provide a fully-informed baseline. In all three conditions the content of the public and private items were counterbalanced, so that each item of information was public in half of all sessions and private in half of all sessions. Table 2 shows these conditions.

[Table 2 about here.]

After subjects read the background information and completed the comprehension quiz, the groups discussed for up to 20 minutes and then chose an option by majority rule. Discussion began with the experimenter announcing each subjects’ leaning to the whole group, After the vote, the experimenter determined the previously unknown state of the world by

¹¹This satisfies the models’s assumption about the relative strengths of the common prior and the private signal.

¹²Even though in the Known-by-All condition the public and private items of information were indistinguishable in subjects’ materials, prior to the begging of each session we designated one item as public and one as private to enable comparisons across condition.

rolling a ten-sided die. If the private and public items favored different options (i.e. each indicated that a different option was more likely to be match), then each possible state of the world had a 50% chance of being chosen. If the private and public items favored the same option, then the state of the world that matched that option had an 80% chance of being chosen. Thus, while the state of the world, and thus the better option, was determined in a manner that subjects knew was stochastic, they knew that the probability that each possible state of the world was related to the information subjects had during discussion.

This experiment examines two dependent variables. First, we recorded whether the group’s decision matched the option that was favored by the private information. Second, we measured the change in subjects’ beliefs caused by discussion. Subjects completed pre- and post-deliberation questionnaires immediately before discussion and immediately after discussion but before the state of the world was determined. The questionnaires included a question asking subjects to predict the percentage chance that each of the two options would turn out to match the state of the world. The wording of the question, altered to fit each decision task, was “At this point you do not know which option will turn out to be better. What percent chance do you think each of the two options has of matching the state of the world?”¹³ The change in these predictions provides an individual-level measure of belief change, and thus a way to measure the influence of the private information in discussion.¹⁴

Two aspects of the experimental design reduce the possibility that an observed effect is the product of the kind of psychological mechanisms described in section 1.2. To reduce the chance that subjects would view majority-minority status as a correctness heuristic, the experimental instructions stressed that there was no relationship between subjects’ leaning and the unknown state of the world.¹⁵ To check subjects’ understanding of this fact a

¹³For example, in the candidate endorsement task described below the question read “What percent chance do you think each of the candidates have of winning the election?”

¹⁴To incentivize participants to give accurate answers, four of the sessions used a proper scoring rule. This was found to have no effect on subject behavior, so the results for these sessions are not reported separately from the other sessions.

¹⁵The instructions stated that “The leanings will have no relationship with the unknown fact,” “the fact that you lean towards one action does not mean that it is more likely to be the better policy, it simply means that all things being equal you prefer that action more than participants who lean in the opposite direction,” and in illustrating the structure of the experiment using the candidate endorsement scenario reiterated that

question in the pre-discussion comprehension quiz asked whether the fact that a subject leaned towards a particular option made that option more likely to be better. To reduce the chance that conformity pressure led subjects to favor the majority-favored option, we reply on private self-reports of beliefs in our analyses of individual belief change. These steps should ensure that any effect seen is the result of the mechanism described in the theoretical model and not alternative mechanisms described by the psychological literature on majority-minority influence literature.

Each session consisted of three rounds. In each round subjects were placed in a new three-person group and were faced with a new decision task that the subjects had not faced before. Rounds were conducted with absolute stranger matching, so that no two subjects were ever in the same group together more than once. Prior to each session the order of the decision tasks, the condition assigned to each round, and the option favored by the public and private items of information were randomly determined. After completing all three rounds, subjects were paid for one of the rounds, which was randomly selected after all three were complete. Average earnings were \$23.97, including a \$15 show-up fee. Each session was conducted by the author as well as two research assistants who were unaware of the study's hypotheses. Subjects were undergraduate students at an elite Northeastern research university who were recruited through the subject pool of a non-deception lab; covariate balance is reported in Appendix B.

3.1.1 Decision Tasks

In each round, groups in the experiment faced one of three decision tasks. Each task consisted of a problem that the group faced, two policy options, and information about the problem and the two options. In the first task, subjects were told that they were part of a political club that must endorse one of two candidates in an upcoming election. The unknown state of the world was which candidate would go on to win the election; a group's decision was matching if the endorsed candidate went on to win. In the second task, subjects represented

“the fact that you lean towards one of the candidates does not mean that candidate is more likely to win the election.” See Appendix C for complete instructions.

a town that must decide whether to expand its sewage treatment facility to accommodate a proposed new factory. Here the unknown state of the world was whether the factory would actually be built, and the matching choice was to expand the treatment facility if the factory was going to be built and not expand the facility if the factory was not going to be built. In the third task the group represented a town that must decide between two plowing contractors, one expensive and one inexpensive, in anticipation of either heavy or light snowfall in the coming winter. In the case of heavy snowfall the expensive contractor would be the matching choice; if snowfall was light, then the cheaper contractor matched the state of the world. Complete descriptions of these tasks are available in online Appendix C.

Figure 1 shows an example of the background information for the candidate endorsement task. The private and public items, bolded here for emphasis (though not in materials given to subjects), are interspersed in the neutral biographies of the candidates. Here, the items are opposed. The public item, which states which candidate has raised more money so far, favors Tom Davies while the private item of information, which states who is going to receive the endorsement of the area's primary newspaper, favors Mark Jones. Both items include some information about the predictiveness of these items of information, ensuring that all subjects know that this information is strongly but not perfectly predictive.

Figure 1 shows the information as it would appear in the background information of an expert subject. In the background information of non-experts these statements were replaced with neutral statements indicating the information was unknown, but that if it were known it would have a particular strength. In the case shown here, the background information given to non-experts would read:

The endorsement of the Springfield Shopper, the highest circulation daily newspaper that reaches about three quarters of the county's population, has not yet been decided. The candidate receiving this endorsement has won in two of the last three elections for county commissioner.

In the Known-by-All condition, all deliberators saw both items of information, with the private item changed slightly so that it was clear that it was known by all group members.

[Figure 1 about here.]

3.2 Results

This paper's hypotheses make predictions about the influence of a private item of information. This experiment measures influence in two different ways. First, information is influential if it is incorporated into deliberators' post-discussion beliefs. To measure this, we examine the change in subjects' beliefs that the policy option favored by the private information is better. Second, information is influential if the decision that it indicates is better is chosen more frequently. These two metrics of influence correspond to this paper's two hypotheses: *H1*, that private information known by a member of the interest minority will have less influence over individual attitudes, and *H2*, that groups will be less likely to incorporate private information known by a member of the interest minority into their decisions.

The first analysis compares change in subjects' beliefs across experimental conditions. We measure the influence of the private information by taking the difference between the pre- and post-deliberation probabilities assigned to the option favored by the private information. We then take the average of these probabilities across non-expert subjects in each experimental condition to test the effects of the experimental manipulation. If information is influential we would expect beliefs to increase; that is, subjects would believe that the option favored by the private information was more likely to be better. If *H1* is correct we would expect beliefs to change less in the Expert-in-Minority condition than in the Expert-in-Majority condition.

The results in Table 3 suggest strong support for this hypothesis. Each row shows the mean belief that the option favored by the private information is better before discussion, after discussion, and the difference between the two. The results show a roughly ten percentage

point change in belief towards the option favored by the private item when the private item is known by a member of the interest majority, significantly greater than the 4.4 percentage point change when the private item is known by a member of the interest minority ($p = .016$, two-sided t-test). Post-deliberation beliefs in the Known-by-All condition, which provides a fully informed baseline, are nearly identical to those in the Expert-in-Majority condition ($p = .688$), but are four percentage points higher than beliefs in the Expert-in-Minority condition, though the relatively small number of subjects in the Known-by-All condition means that this difference is not statistically significant ($p = .166$). In summary, it appears that information known by a member of the minority is about half as influential as the same information known by a member of the majority.

[Table 3 about here.]

Notably, while the change in belief differed across experimental conditions, these changes were also different from the point predictions of the theoretical model. If, as predicted in the theoretical model, private information has *no* influence when the expert is in the minority we would expect no change in beliefs in the Expert-in-Minority condition. The observed change of 4.4 percentage points is significantly different from 0 ($p = .012$). On the other hand, if private information is fully incorporated into non-expert's post-deliberation beliefs when the expert is in the majority we would expect an average change of 15 percentage points in the Expert-in-Majority condition.¹⁶ The observed change of 10 percentage points is significantly different from 15 ($p < .000$). This is not entirely surprising, as existing studies of strategic information transmission generally find support for these models' comparative static predictions but not point predictions (Dickhaut, McCabe, and Mukherji 1995, Cai and Wang 2006). However, it suggests that either embedding the model's signals into explicitly political frames or embedding the exchange of information into face-to-face discussion may weaken, though not eliminate, the effects of these incentives.

¹⁶That is, from a prior belief that the option favored by the public information has a 66.7 percent chance of being correct to a belief that this option has an 80 percent chance of being correct (if the private information also favors this option) or a belief that this option has a 50 percent chance of being correct (if the private information favors the other option).

An additional implication of the theoretical model is that the influence of private information should depend on the majority/minority status of the expert, not the majority/minority status of the non-expert. In other words, experts in the majority should be equally persuasive to non-experts who are in the minority as they are to non-experts who are in the majority. We find evidence that this is the case.¹⁷ Experts in the majority cause a 10.6 percent shift towards the option favored by the private information among other members of the majority, while they cause a 9.1 percent shift among members of the minority; the difference between these figures is not statistically significant ($p = .62$). Of 45 groups in the Expert-in-Majority condition, 21 had a greater change in the non-expert majority member's beliefs while 19 had a greater change in the non-expert minority member's beliefs, and 5 saw both non-experts' beliefs change the same degree. This offers support for the claim that influence does not depend on whether a subject shared the expert's interests, but instead on whether the expert was in the majority.

[Table 4 about here.]

To account for the fact that subjects are embedded in groups, Table 4 reports models that cluster standard errors at the group level.¹⁸ Model one compares subjects in the Expert-in-Majority condition to those in the Expert-in-Minority condition with clustered standard errors; this replicates the comparison in Table 3 and finds the result to still be statistically and substantively significant ($p = .030$). Model Two includes dummy variables for other elements of the groups' decision task: Whether the private signal favored the option that the expert leaned towards, the round number, and the specific task assigned to the group. Model Three controls for the race and gender of the subject as well as the race and gender of the expert in the subject's group.¹⁹ In all of these models the effect of the treatment variable is substantively and statistically significant.

¹⁷We can not analyze whether experts in the minority were equally influential among non-experts as Expert-in-Minority groups had no non-expert members of the minority.

¹⁸A series of Hierarchical Linear models where the experimental condition was a group level predictor produce substantively similar results.

¹⁹None of these demographic controls are statistically significant

[Table 5 about here.]

To test whether this bias affects groups' decisions we examine whether groups made different decisions in the Expert-in-Minority than in the Expert-in-Majority and Known-by-All conditions, with the latter providing a benchmark for fully-informed decision making. Table 5 shows the percentage of groups in each condition that chose the decision that the private item of information indicated was better. In the Expert-in-Majority and Known-by-All conditions groups chose the option favored by the private information 91 and 82 percent of the time; a difference of proportions test shows that the difference between these groups is not statistically significant. In contrast, groups in the Expert-in-Minority condition chose the option favored by the private information a bit less than $\frac{2}{3}$ of the time. This is 29 and 20 percent less often than the rate in the other groups, a substantively significant difference. The difference between the first Expert-in-Majority and Expert-in-Minority conditions is highly significant ($p = .002$, one-sided test), though the small number of groups in the Known-by-All condition mean that the difference between the Expert-in-Minority and Known-by-All is only marginally statistically significant ($p = .099$, one-sided test), though still quite substantively meaningful.

3.3 Discussion

The experimental results provide strong support for the game-theoretic model described in Section 2. The experiment finds that private information known by a member of the minority has less of an impact on other group members' views than the same information known by a member of the majority. Additionally, groups incorporated private information into their decision less often when that information was known by a member of the minority than when it was known by a member of the majority, producing less well-informed decisions.

These results build in important ways on existing experiments that test the effect of interests on communication (Dickhaut, McCabe, and Mukherji 1995, Lupia and McCubbins 1998, Cai and Wang 2006). Other studies have found that people respond to incentives when they communicate. However, "communication" in these studies takes the form of either

sending numerical messages or sending messages over a computer network, much simpler communication methods than the face-to-face discussion in this study. The decision tasks in existing studies are also far more abstract than most political decisions. Scholars may be skeptical of tests of deliberation as strategic information transmission that use highly stylized communication technologies or highly simplified decision environments. Our tests take place in an environment that is more realistic on both counts and finds substantively similar results. The fact that we find support for our directional hypothesis, though not for the model's point predictions, suggests that factors like a more realistic framing of information and face-to-face discussion may introduce incentives, such as a cost to lying, that mitigate the problem identified in Strategic Information Transmission models. However, these factors do not appear to eliminate the basic incentive problem.²⁰

4 Study 2

Study 1 provides clear evidence in favor of the hypothesis that interest minorities are at a disadvantage in small group deliberation. However, the way this study operationalizes two theoretically important elements, the information subjects know and the interests they have, may raise concerns about its external validity. Subjects in the laboratory experiment had interests in the outcome of deliberation that were induced by cash payments, which may differ from normal political interests in important ways. Further, subjects in Study 1 learned everything they knew about the topic under debate from a one-paged vignette that they read immediately prior to discussion. These concerns are in addition to the normal concerns about laboratory experiments conducted with undergraduate subjects.

This section describes a field experiment that attempts to address these limitations. The experiment took place in February, 2011 at a public forum about budget priorities in a school district in New Jersey. As a result of the economic crisis and reduced levels of state aid, the district believed that it might have to make millions of dollars in cuts in the next year's budget. Administrators were eager to gather public input about what cuts should be made

²⁰Thanks to an anonymous reviewer for emphasizing this point.

in order to minimize their impact and provide legitimacy for the cuts. The second factor was particularly important because the district's budget required voter approval. In exchange for conducting a deliberative public meeting and analyzing the data it produced, the district allowed me to control the placement of attendees into groups and the contents of briefing packets attendees received, as well as conduct pre- and post-deliberation surveys.

Residents who attended the meeting completed a pre-survey and then were placed into groups of five, which discussed seven areas where the school district was considering cutting costs or seeking increased revenues. Each deliberator was provided a briefing packet with a small amount of information about each area that their group was to discuss and either one or two of the "sample argument pairs" related to change in that area. We experimentally manipulated the contents of the briefing materials so that one deliberator - the expert - received a packet with an extra sample argument pair related to the area of "sports teams." We assigned subjects to discussion groups so that this expert was in the majority or in the minority in terms of his or her interest in "sports teams." These manipulations allowed me to recreate Study 1's "Expert-in-Minority" and "Expert-in-Majority" conditions in a more realistic environment. When groups finished discussion they voted by majority rule to rank the areas from one to seven, where one was the area where cuts should be made first and seven the area where cuts should be made last. After discussion finished, each deliberator completed a post-discussion survey, including questions measuring the influence of the public and private sample argument pairs.

4.1 Experimental Design

The experiment focused on one of the areas under discussion at the forum: whether or not to cut spending on sports teams, an issue where subjects were expected to have salient and conflicting interests. Cuts to funding in sports teams would likely lead to the elimination of some sports, as well as the imposition of "pay-to-play" fees on students who participated in the remaining sports. These changes would have a large impact on parents of children who played on sports teams (or were expected to in the future), but little or know impact

on other residents,²¹ creating a strong interest on the part of the former group that was not shared by the latter group.

Subjects were classified as pro- or con-sports teams spending cuts based on a ranking task included in the pre-discussion survey. Subjects were presented with a list of seven areas where the school district could reduce costs and ranked these areas from one to seven. Subjects who ranked “sports teams” first through fourth were categorized as “pro,” those who ranked sports teams fifth through seventh were categorized as “con.”²²

Testing the theory required control over two aspects of group discussion: the distribution of information within each group, in order to create an “expert” and the distribution of interests within each group in order to randomly assign the group to the Expert-in-Majority or the Expert-in-Minority condition. Since the interests in this study were naturally occurring we could not randomly assign interests to subjects. Instead, we randomly assigned subjects to groups where the group’s expert was in the majority or minority. Prior to the meeting, we generated a list of groups and randomly assigned each group to an experimental condition. The position (pro or con) that was in the majority was balanced across conditions. The position of each seat in the group was then assigned in order to achieve the balance dictated by the expert’s position and the experimental condition.²³ The result was an ordered list of groups, with a specification of the characteristics of the subject who should sit in each seat in each group.

When subjects completed their pre-meeting survey a research assistant who was blind to the study’s hypotheses assigned them to sit in the first open seat that matched their position, and handed them a briefing packet. Since the order of tables was random, the assignments of subjects to conditions was also random. One group’s data was dropped for failure to comply with the seating assignment. After the meeting began, but before group discussion began,

²¹Unlike in some school districts, high school sports did not play an important role in this communal life of this area, limiting the impact of cuts to those directly affected by them.

²²This definition of “pro” and “con” is admittedly somewhat arbitrary. Tests examining the robustness of these findings to other definitions of “pro” and “con” are available in online Appendix G.

²³For example, if the group was “Expert-in-Majority” and the expert was “pro”, two other seats were assigned to be taken by subjects with a “pro” position and two by subjects with a “con” position.

subjects who were sitting at tables with fewer than five people were reassigned to create as many groups of five as possible. In the interests of mundane realism, the experimenter did not inform subjects of the interests of the other people at their table; instead, subjects were instructed to open discussion by sharing their personal ranking of the seven areas under discussion. This ensured that subjects' interests were public at the start of deliberation.²⁴

After arriving at their group, subjects were encouraged to read their briefing packet, which contained the informational aspect of the experimental manipulation. This three page document contained the list of areas that the group was assigned to discuss, with sample arguments in favor of and against cuts in each area. The argument pairs about sports teams were the following:

- “Other organizations in the community already offer the ability to play some sports, and students could participate there,” and “Reducing the number of sports played in the district would reduce the ability of students to choose a sport that really interests them;”
- “This change does not directly affect classroom instruction, which should be the focus of the district,” and “Reducing the number of students who get to participate in sports would hurt their ability to develop healthy fitness habits and self-esteem.”

These arguments were chosen based on interviews with district staff and members of Parent Teacher Organizations to highlight considerations about the areas that subjects would find to be important, but that they might not have otherwise considered.²⁵ At the request of the district and the Institutional Review Board the arguments were balanced, with one pro and one con in each pair, to reduce the chance that they would have an effect on the post-deliberation preferences of the forum's attendees. In each group, four subjects in the group received one of these argument pairs, with all in the group receiving the same argument pair (henceforth the “public argument pair”). The subject assigned to be the “expert”

²⁴The setting of the public forum meant that we were unable to reliably confirm compliance with this instruction, though anecdotal observations from the research assistants conducting the event suggests that deliberators did, in fact, begin discussion with this information. See section 4.3 for a discussion of the implications of possible non-compliance.

²⁵We also attempted to choose arguments that appeared to be of roughly equal importance, so as to satisfy the requirements of the theoretical model.

received a briefing packet with both argument pairs (including what we henceforth refer to as the “private argument pair”). Which of the two argument pairs was public was balanced across experimental conditions. The manipulation, varying the appearance of two sample arguments in a briefing packet in advance of a two-hour deliberation, is intentionally weak, in order to avoid exerting too much influence over the results of a public forum and to provide a hard test for the theory.

Since the arguments were balanced, we expect the argument pairs to no net effect on subjects’ post-deliberation attitudes. This means we cannot measure the influence of each argument pair using change in deliberators’ opinions, as we did in Study 1, and must instead rely on subjects’ self-reported reactions to these arguments. We do so using three items on the post-discussion survey:

Some people say that [Argument Text]

1. How much thought have you given this claim tonight? 1. A great deal of thought 2. Some Thought 3. Not much thought 4. No thought/Never heard this claim
2. How important is this claim to deciding how to rank the area of sports teams? 1. Very important 2. Somewhat important 3. Somewhat unimportant 4. Not important at all
3. Please indicate how much you agree or disagree with this claim. 1. Strongly agree 2. Somewhat agree 3. Neither agree nor disagree 4. Somewhat disagree 5. Strongly disagree

These items were repeated for each of the four arguments about sports teams.

80 people attended the forum; covariate balance is reported in Appendix B. Groups discussed on average for 113 minutes.

4.2 Results

As with the first study, analysis of data from this study will focus on the influence of arguments on non-experts, those subjects who did not have the private argument pair in their briefing packets. Fifteen groups from the study provided usable data. The expert in each group is excluded; further, four non-expert subjects in these groups did not complete

the post-survey and were thus dropped from the data. The result was 56 observations, 23 in the Expert-in-Minority condition and 33 in the Expert-in-Majority condition.

We measure the influence of the public and private argument pairs by taking the average of the response to the two arguments in the argument pair on each of the three post-survey questions measuring influence. Since this measure of influence is the average influence of two arguments favoring different positions it is not affected by subjects' tendency to find arguments compatible with their position more convincing.²⁶ We test whether arguments made by members of the minority were less influential than the same arguments made by members of the majority by comparing reactions to the arguments in the private argument pair across the two experimental conditions.

[Figure 2 about here.]

The results confirm the findings from Study 1 that arguments made by members of the interest minority are less influential. Figure 2 shows the average reaction of non-expert group members to the arguments in the private argument pair after discussion was over, along with 90 percent confidence intervals. Only the measure of agreement shows a statistically significant difference between the two ($p=.020$, one-sided t-test). On average, subjects rated their agreement with arguments known by experts who were in the minority .54 points less on a five point scale, or 60 percent of a standard deviation. This finding suggests that arguments known by members of the interest minority are not given less thought, nor are they deemed less important because they are introduced by members of the minority. Instead, deliberators are simply more likely to disagree with arguments made by a member of the interest minority than they are to disagree with the same argument when it is introduced by a member of the interest majority.

An additional implication of the model is that experts in the majority should be credible to all non-experts, including members of the minority, while experts in the minority should not be credible to any non-experts, including other members of the minority. To test this,

²⁶See Online Appendix F for an auxiliary analysis of individual arguments.

Figures 3 and 4 replicate Figure 2 looking at self-reported influence only among members of the majority and minority, respectively. Members of the interest majority and members of the interest minority appear to be similarly distrustful of arguments known by members of the interest minority. These differences are of similar magnitude to those shown in Figure 2, though the small sample size means that this difference is not significant by the traditional $p=.05$ standard ($p=.053$ and $.108$ respectively, one-sided t-test). Still, it seems that members of the interest minority have a hard time making their case to anyone, not just members of the interest majority. This finding is consistent with the predictions of the theoretical model; the incentive a member of the interest minority may have to misrepresent information means that his or her statements are less credible to everyone, including other members of the interest minority.

[Figure 3 about here.]

[Figure 4 about here.]

[Table 6 about here.]

Table 6 shows the results of bivariate models with clustered standard errors that replicate the finding that is shown in Figures 2, 3, and 4. As in the analysis in Study One, the treatment is represented by a dummy variable that equals one if the group is in the “Expert-in-Majority” condition. Given the small number of observations in the experiment, we did not attempt to run these models including covariates. For each dependent variable the first row shows the results using all subjects, while the subsequent rows show the results using minority subjects only and majority subjects only. As suggested by Figure 2, being in the Expert-in-Majority condition caused a statistically significant increase in agreement with the private argument pair but no significant effect on the “thought about” or “considered important” variables. The results from the majority-only and minority-only models largely replicated these findings, though once again the sample size for these models is quite small.

4.3 Discussion

The results from this study replicate the findings of Study 1: Arguments known by members of the interest minority are less influential than the same arguments known by members of the interest majority. Since the content of these arguments are the same in each condition, and the majority-minority status of the deliberators who knew them was randomly varied, we can conclude that it is the interest majority-minority status of the person who knows the arguments, and not some other factor, that is producing these results. Interestingly, deliberators do not appear to have thought less about or consider less important arguments known by members of the interest minority. Instead, they are less likely to believe these arguments when they are known by a member of the interest minority.

The field setting for this study requires giving up some control over the key moving parts in the theory: interests and information. In study 1 we could be sure that subjects stated their interests publicly at the start of the session; in the field, we depended on subjects to comply with instructions to state their preferences at the beginning of discussion. If compliance was not perfect this would mark a deviation from the model described in section 2 in that subjects' interests would be private instead of public information. Meirowitz (2007) argues that in such situations fully revealing equilibrium are more common, though this depends on the assumption that subjects over-estimate the probability that others share their own interest. However, if this is the case it would weaken the strength of the treatment, in that some members of the minority would have more credibility than if their interests were publicly known, and thus provide a harder test of *H1*.

In terms of information, we might be concerned that deliberators who were not assigned to know the private argument pair might nevertheless be aware of and raise these arguments in discussion. Indeed, in the interest of external validity the arguments were designed to look like the kinds of arguments that deliberators might bring up themselves. In a sense, the treatment might be thought of as raising the salience of a particular argument with one deliberator, increasing the chance that this argument would be raised by the target deliberator instead of someone else. Again, this deviation from the theoretical model should weaken

the strength of the treatment, providing a harder test of *H1*. These possible deviations do suggest some caution in interpreting the null findings of effects of the manipulation on the “thought about” and “consider important” dependent variables.

The field setting also makes it somewhat harder to rule out alternative psychological mechanisms of influence. As in Study 1, the use of private survey responses as the dependent variable reduces the effect of conformity pressures. Additionally, the fact that groups were divided three/two meant that minorities members were not alone in their minority status, further reducing conformity pressures. Unlike in Study 1, in this Study we could not use strong experimental instructions to clearly state that there was no connection between subjects’ private interests and the objective correctness of cutting or not cutting sports team funding. Instead, we rely on fact that cutting sports teams would primarily impose a private cost on a subset of forum participants - those who had or expected to have students who played on school sports teams - to reduce the possibility that majority status would exert influence through a correctness heuristic.

Finally, the field setting of this experiment also places limitations on sample size and measurement. The number of subjects in the field experiment was limited by the number of residents of a small New Jersey town who were willing to discuss school board budgets for two hours; as a result, the study is underpowered in some places. Further, the ethical requirement to avoid influencing deliberators’ policy preferences meant that we could not use the change in deliberators’ opinions or the change in groups’ decisions as measures of an argument’s influence. Instead, we rely on deliberators’ self-reports of how they responded to each argument. This dependent variable also makes it hard to compare the magnitude of effects across the two studies. Nevertheless, the field environment offers naturally occurring interests and a more natural information environment than Study 1. In short, Study 2 is weak in areas where Study 1 is strong, and vice versa; that both suggest a similar conclusion should give greater confidence than either study could by itself.

5 Conclusion

The “deliberative turn” in democratic theory (Dryzek 2000) has inspired an extensive body of empirical work that has shown the promise of deliberation, but also demonstrated a number of problems or biases in small-group deliberation. Deliberation can be biased against women (Karpowitz and Mendelberg 2014); deliberation can produce opinion polarization (Schkade, Sunstein, and Hastie 2007); deliberation with too much or too little preference heterogeneity can be low quality (Esterling, Fung, and Lee 2015). This paper adds another bias to this list: a bias against information introduced by members of the interest minority. This bias may in some ways be more intractable than racial or gender bias; while we might hope that racial and gender bias in society will eventually be reduced or eliminated, “irresolvable conflicts in opinions and interests” are an inherent feature of living in a pluralistic society (Mansbridge et al. 2010, pg. 64-65).

Some choices in the modeling and experimental design presented in this paper place important limitations on the generalizability of these results to all contexts. The model assumes that deliberators’ interests are publicly known, though this may not be the case in all deliberative situations (cf. Meirowitz 2007). Further, while the theoretical model accommodates a wide range of types of interests, the studies in this paper deal primarily with self-interest. This was done for methodological reasons; self-interest is relatively easy to measure as well as to manipulate in the laboratory. While we believe that these results have relevance for the impact of interests defined broadly, future work should examine whether different kinds of interests (e.g. moral values vs. ideological beliefs) have different effects on the incentive problem described here.

While (with these caveats) these results pose a challenge for deliberative democrats, they do not in any way invalidate democratic deliberation or the normative theories underlying it. Instead, we hope that they will lead deliberative democrats to consider ways to ameliorate the incentive problem described here (cf. Landa and Meirowitz 2009). Such strategies fall into two broad categories: reducing the importance of private information to deliberation

(generally by making this information publicly known), or reducing the incentive to misrepresent private information in deliberation. The most common example of the former is providing deliberators with information from experts who are perceived to be neutral, as is done in many mini-public designs such as Citizens' Juries or Deliberative Polls (Fishkin and Luskin 2005).²⁷ While valuable, this strategy is limited in important ways. They are costly, both in terms of money and in terms of deliberators' time. Further, they are limited to types of information where neutral experts are available, such as uncontroversial technical information. On some issues, the politicization of information may make it impossible to find experts that are perceived to be neutral by all deliberations. Further, using experts may be impossible for many important kinds of information, such as moral arguments or situated knowledge that is known only by deliberators themselves (Young 1996).

Alternatively, deliberative forums can attempt to make misrepresenting information costly, either through institutions or norms, and in doing so make all speech by interest minorities more credible. Indeed, some advocates of deliberation argue that the incentive problems described in formal models are not a problem for deliberation because norms of truth-telling will prove far stronger than any incentive to misrepresent information (Steiner 2008, Neblo 2015). We are not so optimistic as to believe that norms of truth-telling will be sufficiently strong in all deliberative situations to eliminate this incentive problem, and believe that the studies presented in this paper are evidence of this.²⁸ However, where such norms exist they may be able to ameliorate the incentive problem described in this paper; understanding the effect of norms on this incentive problem is an important task for future research. Also important is studying how and when such norms develop as well as the effect of things like education

²⁷An alternative strategy used by participatory budgeting programs takes deliberators on a tour of affected neighborhoods, making information about the needs of different communities public (Wampler 2007).

²⁸It is notable that the most prominent attempt to develop a holistic measure of deliberative quality explicitly disavows attempting to measure truthfulness or sincerity of speech (Steenbergen, Bächtiger, Spöndlin, and Steiner 1995, pg. 26), suggesting that empirical evidence in support of this argument may be hard to come by. An alternative version of this argument is that situations where citizens allow their self-interest to affect the exchange of information should not be counted as instances of "deliberation." While we do not wish to embark on a debate about what does or does not count as "real deliberation," (for such a debate see Steiner (2008), Schneider (2008), and Goodin (2008)) we suggest that the results presented in this article can be interpreted as helping to delineate situations where deliberation as defined by authors like Steiner and Neblo is more or less likely to be possible.

or trust-building dialogue on their development. In places where norms are not sufficiently strong to solve this problem, intuitional features might be used to introduce external incentives for truth-telling (cf. Meirowitz 2006).

Still, in cases where there is a great deal of intractable interest conflict, deliberation may not be the answer. Instead, society may have to revert to aggregative models of democracy, which protect the equality of citizens through equal vote instead of equal voice. This argument is not new; indeed, it mirrors Mansbridge's (1980) argument about the difference between unitary and adversarial democracy. Mansbridge (1980) argues that voice, by itself, is only appropriate for unitary democracy, or democratic communities where all share the same interests. When interests are in conflict, discussion will marginalize socially dominated groups and dissident citizens as those in the majority attempt to present their interests as the common will of the community. In such situations the secret ballot, the most important feature of adversarial democracy, is necessary to protect the equality of all and voice or discussion is less important for democratic legitimacy. The mechanism underlying the marginalization of interest minorities that is described here is different than in Mansbridge's account. However, the caution it recommends about deliberation in situations of strong, conflicting interests is the same.

6 Acknowledgements

The author thanks Tali Mendelberg, Adam Meirowitz, Martin Gilens, Larry Bartels, Chris Karpowitz, Shawn Rosenberg, Chuck Myers, Michael Neblo, Joshua Pasek, Dave Glick, Dustin Tingley, Ben Lauderdale, Matthew Wilson, participants in Princeton's American Politics Research Seminar, the University of Michigan's Interdisciplinary Workshop on American Politics and the New York Area Political Psychology working group, as well as several anonymous reviewers for helpful feedback, discussion, and comments on this project. Many thanks also to Jason Anton, Elizabeth Ingriselli, Anand Krishnamurthy, Nora Xu, Fiona Wilson, Mark Benjamin, Shaina Watrous, Kai Khor, and Pritha Dasgupta for research assistance.

All errors, be they of omission or commission, are of course my own.

References

- Austen-Smith, David, and 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, and James Dugan. New York: Springer.
- Austen-Smith, David, and Timothy J. Feddersen. 2006. "Deliberation, Preference Uncertainty, and Voting Rules." *American Journal of Political Science* 100(2): 209–217.
- Cai, Hongban, and 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, and Phillip Petit. Boston: Blackwell pp. 17–34.
- Crawford, Vincent, and Joel Sobel. 1982. "Strategic Information Transmission." *Econometrica* 50: 1431–1451.
- Dickhaut, John W., Kevin A. McCabe, and Arijit Mukherji. 1995. "An experimental study of strategic information transmission." *Economic Theory* 6(3): 389–403.
- Dickson, Eric S., Catherine Hafer, and Dimitri Landa. 2008. "Cognition and Strategy: A Deliberation Experiment." *Journal of Politics* 70(4): 974–989.
- Dryzek, John S. 2000. *Deliberative Democracy and Beyond: Liberals, Critics, Contestations*. New York: Oxford University Press.
- Esterling, Kevin. 2011. "'Deliberative Disagreement' in U.S. Health Policy Committee Hearings." *Legislative Studies Quarterly* 36(2): 169–198.
- Esterling, Kevin M., Archon Fung, and Taeku Lee. 2015. "How Much Disagreement Is Good for Democratic Deliberation? The California Speaks Health care Reform Experiment." *Political Communication* 32(4): 529–551.
- Estlund, David. 1997. "Beyond Fairness and Deliberation: The Epistemic Dimension of Deliberative Authority." In *Deliberative Democracy: Essays on Reason and Politics*, ed. James Bohman, and William Rehg. Cambridge: MIT Press pp. 173–204.
- Farrell, Joseph, and Matthew Rubin. 1996. "Cheap Talk." *Journal of Economic Perspectives* 10(3): 103–118.
- Fishkin, James S., and Robert C. Luskin. 2005. "Experimenting with a Democratic Ideal: Deliberative Polling and Public Opinion." *Acta Politica* 40: 284–298.
- Gerardi, Dino, and Leeat Yariv. 2007. "Deliberative Voting." *Journal of Economic Theory* 134(1): 317–338.
- Goodin, Robert E. 2008. "Deliberative Lies." *European Political Science* 7(2): 194–198.

- Goodin, Robert E. 2012. *Innovating Democracy: Democratic Theory and Practice After the Deliberative Turn*. Oxford: Oxford University Press.
- Grogan, Colleen M., and Michael K. Gusmano. 1997. *Healthy Voices, Unhealthy Silence: Advocacy and Health Policy for the Poor*. Washington: Georgetown University Press.
- Gutmann, Amy, and Dennis Thompson. 1996. *Democracy and Disagreement*. Cambridge: Cambridge University Press.
- Habermas, Juergen. 1975. *Legitimation Crisis*. Boston: Beacon Press.
- Karpowitz, Christopher F., and Tali Mendelberg. 2014. *The Silent Sex: Gender, Deliberation, and Institutions*. Princeton, NJ: Princeton University Press.
- Knight, Jack, and James Johnson. 1997. "What Sort of Equality does Deliberative Democracy Require?" In *Deliberative Democracy: Essays on Reason and Politics*, ed. James Bohman, and William Rehg. Cambridge: MIT Press pp. 279–320.
- Landa, Dimitri, and Adam Meirowitz. 2009. "Game Theory, Information, and Deliberative Democracy." *American Journal of Political Science* 53(2): 427–444.
- Lupia, Arthur, and Mathew D. McCubbins. 1998. *The Democratic Dilemma: Can Citizens Learn What They Need to Know?* Cambridge: Cambridge University Press.
- Mackie, Diane M. 1987. "Systematic and Nonsystematic Processing of Majority and Minority Persuasive Communications." *Journal of Personality and Social Psychology* 53: 41–52.
- Mansbridge, Jane J. 1980. *Beyond Adversary Democracy*. New York: Basic Books.
- Mansbridge, Jane, James Bohman, Simone Chambers, David Estlund, Andreas Føllesdal, Archon Fung, Cristina Lafont, Bernard Manin, and José Luis Martí. 2010. "The Place of Self-Interest and the Role of Power in Deliberative Democracy." *Journal of Political Philosophy* 18(1): 64–100.
- Martin, Robert, Miles Hewstone, Pearl Y. Martin, and Antonis Gardikiotis. 2008. "Persuasion from Majority and Minority Groups." In *Attitudes and Attitude Change*, ed. William D. Crano, and Radmila Prislin. New York: Psychology Press pp. 189–212.
- Meirowitz, Adam. 2006. "Designing Institutions to Aggregate Preferences and Information." *Quarterly Journal of Political Science* 1: 373–392.
- Meirowitz, Adam. 2007. "In Defense of Exclusionary Deliberation: Communicating and Voting with Private Beliefs and Values." *Journal of Theoretical Politics* 19(3): 301–327.
- Mendelberg, Tali, and John Oleske. 2000. "Race and Public Deliberation." *Political Communication* 17(2): 169–191.
- Moscovici, Serge. 1980. "Toward a Theory of Conversion Behavior." In *Advances in Experimental and Social Psychology*, ed. L. Berkowitz. New York: Random House pp. 209–239.

- Myers, C. Daniel, and Tali Mendleberg. 2013. "Political Deliberation." In *Oxford Handbook of Political Psychology*, ed. Leonie Huddy, David O. Sears, and Jack S. Levy. New York: Oxford University Press pp. 699–736.
- Neblo, Michael A. 2015. *Deliberative Democracy Between Theory and Practice*. New York: Cambridge University Press.
- Sanders, Lynn. 1997. "Against Deliberation." *Political Theory* 25(3): 347–376.
- Schkade, David, Cass R. Sunstein, and Reid Hastie. 2007. "What Happened on Deliberation Day?" *California Law Review*, 95(3): 915–940.
- Schneider, Gerald. 2008. "De Deliberatibus Disputandum est: A Response to Jürg Steiner." *European Political Science* 7(2): 199–206.
- Stasser, Garold, and William Titus. 2003. "Hidden Profiles: A Brief History." *Psychological Inquiry* 14(3-4): 304–313.
- Steenbergen, Marco R., Andre Bächtiger, Markus Spörndlib, and Jürg Steiner. 1995. "Measuring Political Deliberation: A Discourse Quality Index." *Economic Theory* 6(3): 389–403.
- Steiner, Jürg. 2008. "Concept Stretching: The Case of Deliberation." *European Political Science* 7(2): 186–206.
- Thompson, Dennis. 2008. "Deliberative Democratic Theory and Empirical Political Science." *Annual Review of Political Science* 11: 497–520.
- Walsh, Kathy Cramer. 2007. *Talking About Race: Community Dialogues and the Politics of Difference*. Chicago: University of Chicago Press.
- Wampler, Brian. 2007. *Participatory Budgeting in Brazil: Contestation, Cooperation, and Accountability*. University Park: Pennsylvania State University Press.
- Young, Iris Marion. 1996. "Communication and the Other: Beyond Deliberative Democracy." In *Democracy and Difference: Contesting the Boundries of the Political*, ed. Selya Benhabib. Princeton: Princeton University Press pp. 120–136.

7 Biographical Statement

C. Daniel Myers is an Assistant Professor at the University of Minnesota - Twin Cities, Minneapolis, MN, 55455.

*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. **You have information that the Springfield Shopper, the highest circulation daily newspaper that reaches about three quarters of the county's population, has decided to endorse Jones. The candidate receiving this endorsement has won in two of the last three elections for county commissioner. *The other***

members of the group are not aware of this endorsement. 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.** 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 1: Sample Candidate Biographies

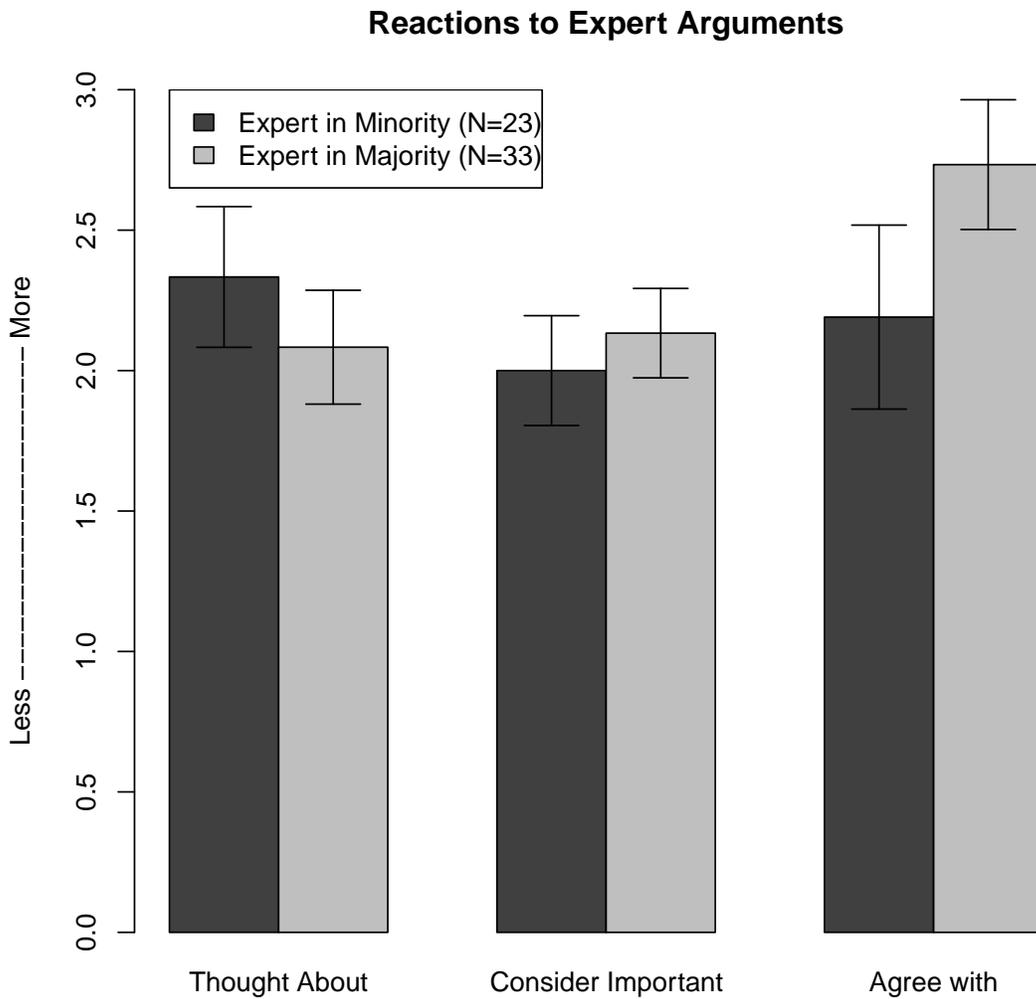


Figure 2: Post-Discussion Reactions to Expert Information by Whether Expert Was in the Majority

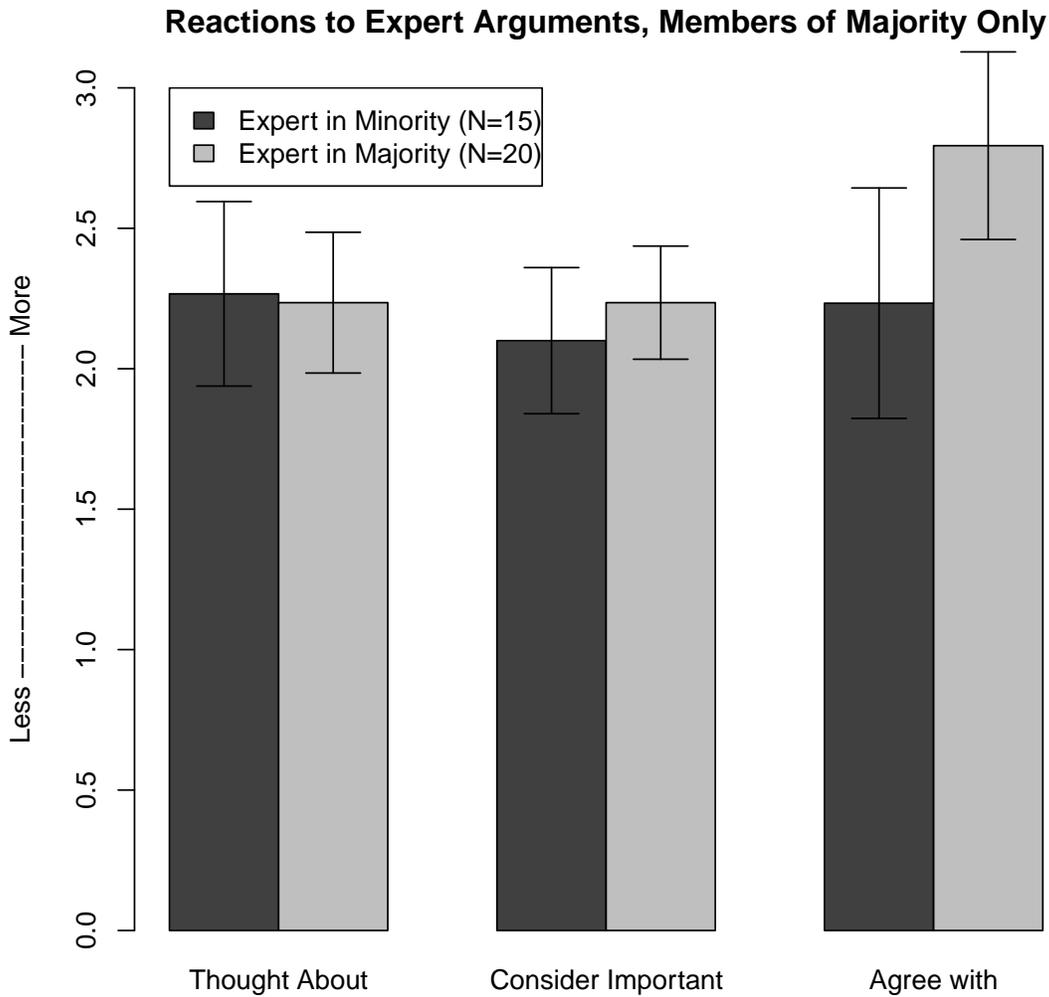


Figure 3: Post-Discussion Reactions to Expert Information Among Members of the Majority

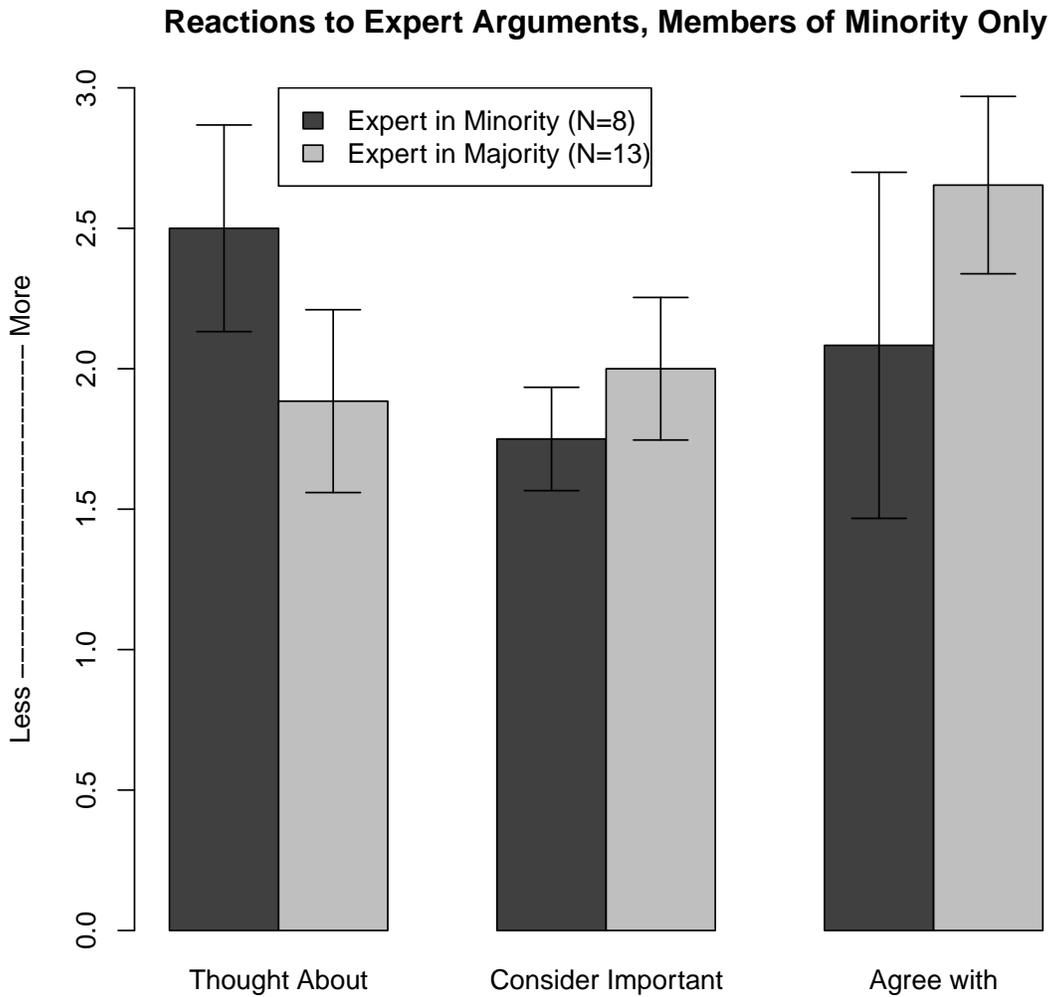


Figure 4: Post-Discussion Reactions to Expert Information Among Members of the Minority

Group Choice	Earnings
Group Choice = State of World and Group Choice = Subject's Leaning	\$15
Group Choice = State of World and Group Choice \neq Subject's Leaning	\$10
Group Choice \neq State of World and Group Choice = Subject's Leaning	\$5
Group Choice \neq State of World and Group Choice \neq Match Subject's Leaning	\$0

Table 1: Possible Payoffs

Condition	Expert-in-Majority			Expert-in-Minority			Known-by-All		
	1(Maj)	2(Maj)	3(Min)	1(Maj)	2(Maj)	3(Min)	1(Maj)	2(Maj)	3(Min)
Subject(Interest)									
Knows Public Info?	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Knows Private Info?	Yes	No	No	No	No	Yes	Yes	Yes	Yes

Table 2: Study One Experimental Conditions

Condition	Before Discussion	After Discussion	Difference
Expert-in-Minority (N=86) (95% Conf Int)	.524 (.494, .555)	.568 (.528, .608)	.044 (.010, .077)
Expert-in-Majority (N=90) (95% Conf Int)	.498 (.465, .531)	.597 (.560, .633)	.099 (.069, .123)
Known-by-All (N=66) (95% Conf Int)	.599 (.567, .631)	.608 (.569, .647)	.008 (-.020, .037)

Table 3: Belief that Option favored by Private Information is Better

	Model 1	Model 2	Model 3
Expert in Majority	.056* (.025)	.062** (.027)	.078** (.028)
Private Signal = Expert Leaning	—	-.020 (.031)	-.018 (.028)
Round 2	—	-.001 (.033)	.002 (.034)
Round 3	—	-.046 (.031)	-.053 (.031)
Task 2	—	-.011 (.039)	-.014 (.035)
Task 3	—	.005 (.028)	.016 (.025)
Constant	.042 (.021)	.067 (.031)	.024 (.041)
Demographic Controls?	N	N	Y
N	176	176	176
*: $p < .05$. **: $p < .01$			

Table 4: Effect of Expert's Majority Status on Belief Change, Clustered Standard Errors

	Expert-in-Minority	Expert-in-Majority	Known-by-All
Percentage	.62	.91	.82
# of Groups	43	45	22

Table 5: Percentage of Groups Choosing Policy Favored by Private Information

	Agree With		
	All Subjects	Minority Only	Majority Only
Expert in Majority	.54** (.14)	.57 (.42)	.56* (.23)
Constant	2.81*** (.08)	2.92*** (.37)	2.77*** (.18)
Multiple R^2	.09	.05	.03
	Thought About		
	All Subjects	Minority Only	Majority Only
Expert in Majority	-.25 (.22)	-.62 (.32)	-.03 (.22)
Constant	1.67*** (.15)	1.50*** (.22)	1.73*** (.14)
Multiple R^2	.03	.16	.00
	Consider Important		
	All Subjects	Minority Only	Majority Only
Expert in Majority	.13 (.13)	.25 (.18)	.14 (.16)
Constant	2.00*** (.08)	1.25*** (.12)	1.90*** (.10)
Multiple R^2	.01	.06	.01
n	55	21	35
$n(\text{clusters})$	15	15	15

*** $p < .001$; ** $p < .01$; * $p < .05$

Table 6: Effect of Expert's Majority Status on Reaction to Arguments, Clustered Standard Errors