Decision-making in Committees

Li, Hao

Department of Economics, University of Toronto

Wing Suen

School of Economics and Finance, University of Hong Kong

February 1, 2009

Abstract: This article reviews recent developments in the theory of committee decision-making. A committee consists of self-interested members that make a public decision by aggregating imperfect information dispersed among them according to a pre-specified decision rule. We focus on costly information acquisition, strategic information aggregation, and rules and processes that enhance the quality of the committee decision. Seeming inefficiencies of the committee decision-making process such as over-cautiousness, voting, and delay emerge as partial remedies to these incentive problems.

JEL classification: D71, D82

Corresponding author: Li, Hao, 150 St. George Street, University of Toronto, Toronto, Ontario M5S 3G7. Email: li.hao@utoronto.ca. Phone: 1-416-978-1970

1. Introduction

Decision-making by committees is the subject of much ridicule. Committees are said to be incompetent (*Committee is a group of men who individually can do nothing but as a group decides that nothing can be done*), to produce monstrous decisions (*A camel is a horse designed by committee*), and to be painfully slow (*A committee keeps minutes and wastes hours*). Yet they are also a ubiquitous fixture of modern society, because the greatest strength of making decision by committees lies in the ability to draw diverse viewpoints from constituent members. This was first articulated more than two centuries ago by Condorcet (1875) in what became known as the Condorcet jury theorem, which shows that voting groups with diverse information make better decisions the larger the group size. The Condorcet jury theorem is a statistical proposition based on an early application of the law of large numbers. Economists have increasingly become interested in this topic, and the economics literature develops by placing incentives at the center.

In our view, an economic model of committee decision-making must capture the following three features. First, the decision to be made by the committee is public in the sense that no constituent member can escape from the consequence of the decision. The public nature of the committee decision distinguishes the committee decision-making framework from models of allocating private consumption goods, such as auctions. Second, the committee makes the decision by aggregating the relevant information dispersed among its members through a pre-committed rule. Analysis of committee decision-making involves different insights from both models of axiomatic social choice theory, and models of strategic information transmission such as signaling and cheap talk. Lastly, committees members are self-interested economic agents that have different preferences over the decision given any information. Understanding how decisions are made in a committee requires us to analyze information aggregation by treating each information source as a strategic agent rather than a statistical unit.

The theoretical foundation of analyzing committee decision-making is the literature on mechanism design with multiple agents. Unfortunately, a large part of this theoretical literature employs monetary transfers to resolve incentive problems when agents have asymmetric information and heterogeneous preferences. Such solutions are rarely relevant in observed practices of committees. It is thus unsurprising that the literature on committee decision-making emerged from different strands of the political science literature, where monetary transfers are ruled out as infeasible. In the literature on agenda-setting, Gilligan and Krehbiel (1989) first introduce asymmetric information and heterogeneous preferences to study the effects of open versus closed amendment rules in legislatures. Austen-Smith (1990) uses a similar framework to address the issue of whether pre-voting debates, modeled as cheap talk communication, can affect legislative decisions. In the literature on elections, Austen-Smith and Banks (1996) first point out that sincere voting as presumed in Condorcet (1875), that is, voting according to one's own private signal only, is neither welfare-maximizing nor consistent with rationality even when committee members have identical preferences. Strategic information aggregation requires instead conditioning one's vote on the information inferred from the vote being pivotal as well as on one's own private information. Implications of this "pivotal voting" argument are explored in a series of influential papers on elections by Feddersen and Pesendorfer (1996; 1997; 1998).

In this article we review some of the recent developments in the literature on committee decision-making, using themes from our own published works and some ongoing research with Ettore Damiano to organize the exposition (Li, 2001; Li, Rosen and Suen, 2001; Li and Suen, 2004; Damiano, Li and Suen, 2008b). We focus on three sets of interrelated incentive issues involved in decision-making in committees. First, even if committees are efficient aggregators of diverse information, what are the incentives for individual committee members to acquire information in the first place? Second, given that members of a committee have different goals and preferences, what are the incentives for them to reveal their private information in the decision making process? Third, when committee members behave strategically, what voting rule and committee composition will best enhance the quality of decisions? A main message we would like to convey is that when these incentive problems are taken into account, some of the seeming inefficiencies of the committee decision-making process can be cast in a more positive light.

In section 2 we introduce a prototype model of a committee which has to make a decision between a status quo and an alternative.[1] The committee has limited information, dispersed among its members, to base its decision on. This means that two kinds of mistakes are possible: adopting the alternative when it is in fact inferior to the status quo (type I error) and rejecting the alternative when it is superior (type II error). Committee members may have conflicting interests either because they have different assessments about how the two choices compare or because they weigh type I error and type II error differently. The model captures some key features common to real-life situations of making decisions in committees. For example, in a corporation boardroom meeting to decide whether or not to approve a merger, there may be sufficient uncertainty about the merits of the merger that the board has to worry about both the possibility of rushing to an unprofitable deal and the possibility of missing a grand opportunity. Board members typically have different interests in a merger: managers may be fearful of losing out to the new management team, while investors are more interested in expanding the business empire. Nevertheless such divergence of interest is not absolute. If the evidence in favor of one of the decisions is sufficiently strong, all members would agree to taking that decisions. Hence the pooling of information is an essential part of the decision making process.

Our model in section 2 assumes that the private information of each member is a continuously distributed signal.[2] The assumption of continuous signals distinguishes the present model from some earlier models of group decision-making (e.g., Austen-Smith and Banks, 1996; and Feddersen and Pesendorfer, 1996, 1997, 1998; Chwe, 1999). It is adopted for two reasons. First, the assumption allows for a richer analysis of strategic manipulation of private information in committees. Second, and more importantly, since any voting procedure cannot completely reflect the strength of evidence contained in continuous signals, making decisions by voting cannot be the first-best way of aggregating information. This raises the important issue of why voting is such a ubiquitous feature of committee

[1] The framework is borrowed from optimal statistical decisions (DeGroot, 1970). The problem of choosing between the status quo and the alternative for given evidence is a standard hypothesis testing problem. Of course, incentive issues in gathering evidence, manipulating information or delegating decisions are not part of the statistics literature.

[2] The model presented here also assumes that the signal is one-dimensional, but this is not an essential feature of our model. Generalization to the multi-dimensional case is straightforward.

decision-making.[3]

Section 3 builds on the prototype model by introducing an information acquisition decision. We begin by noting that a free-rider problem arises when the committee decision has to be made based on privately collected evidence, leading to insufficient effort in gathering evidence. To alleviate the free-rider problem, the committee can commit to a "conservative" rule, whereby the decision is made against the alternative favored by the committee's dominant preference when the evidence supports it but is not preponderant. In the example where all committee members have the same ex ante preference that is biased in favor of the alternative, a conservative rule corresponds to a higher standard of evidence needed to adopt the alternative than what is ex post optimal. Optimal conservatism increases private incentives to gather evidence and improves the quality of the committee decision. This result explains why sometimes committees appear overly cautious toward favored alternatives.

Gersbach (1995) is among the first to recognize the free-rider problem in voting in committees. The contribution here lies in identifying an ex post inefficient rule as a remedy to this problem. It is well-known in the mechanism design literature that committing to ex post suboptimal allocation rules may provide appropriate incentives for agents in moral hazard and asymmetric information problems. More closely related, in the absence of monetary transfers, a principal can provide greater incentives for an agent to gather information through appropriate decision rules, such as restrictions on amendments as in Gilligan and Krehbiel (1987) and delegation as in Aghion and Tirole (1997). As in these papers, the main idea outlined in section 3 is not to achieve the first best through some clever mechanisms without the help of monetary transfers, but to provide an intuitive explanation for the kind of inertia or conservatism often observed, and criticized, in committee decision-making.

Section 4 sets aside the problem of information acquisition and focuses on the strategic issues that arise when committee members have conflicting preferences on the optimal decision given the available information. The degree of conflict between two members

[3] Klevorick, Rothschild, and Winship (1984) demonstrate how the quality of verdict improves through optimal information aggregation, as opposed to a simple majority vote among jurors.

of a committee can be represented by the difference in their ex post optimal standards of evidence. A key result is that voting procedures emerge endogenously in committee decision-making: it coarsens the transmission of each member's information, but controls strategic manipulations and allows some degree of information sharing. This result may be viewed a natural extension of Crawford and Sobel's (1982) work on strategic information transmission from an informed sender to an uninformed receiver. Like the sender in their cheap talk model, in any incentive compatible outcome of transforming the signals of the members to a decision, each member represents his information in discrete categories. In our committee decision-making framework, these categories correspond to votes, and an incentive compatible outcome corresponds to a general voting procedure that aggregates members' votes into a decision. Moreover, similar to Crawford and Sobel (1982), smaller conflicts among members allow for finer partitions of their information, and thus greater sophistication in an incentive compatible voting procedure. Unique to our model of strategic information aggregation, however, are interactions in the way each member partitions his information in categories. A committee member who is more biased toward the status quo will set a lower threshold for voting for the status quo if the other committee member becomes more biased in the opposite direction, but he will also cast the pivotal vote less often if the other committee member becomes better informed. Although voting outcomes cannot be ex post Pareto optimal, voting represents an incentive compatible way of aggregating private information. Moreover the gains from information sharing exceeds the costs from strategic manipulation in the sense that each committee member prefers committee decision making to having one member dictate the outcome on the basis of his own information.

Our model in section 4 is closely related to the literature on strategic voting with discrete data (Austen-Smith and Banks, 1996; Feddersen and Pesendorfer, 1996; 1997; 1998), which identifies "pivotal voting" as the primary mode of strategic interactions among voters. In these models, private signals are discrete, a feature that leads the authors to presume voting procedures as the only means of information aggregation, and limits their analysis of information manipulation to mixed strategies voting and uninformative voting, i.e., voting for the ex ante bias regardless of the private information. By assuming continuous signals in our model, we derive voting as an incentive compatible outcome of information manipulations. In our analysis of strategic voting, we also go beyond the logic of pivotal voting to study a richer set of information manipulations in committee decision-making.[4] In particular, we provide an analysis of two-category voting that emphasizes the role of conflicting interests in committee decision-making, and address how conflicts in the committee affect manipulation and sharing of information.

In section 5, we ask questions about delegation (why and when a decision-maker should delegate the decision to a committee) and recruitment (how the decision-maker should pick committee members). Unlike previous papers that stress the benefits of delegation in terms of providing incentives for committee members to gather information (Gilligan and Krehbiel, 1989; Dewatripont and Tirole, 1999), we focus on the use of delegation mechanisms by principals to mitigate the agency problem of information manipulation. We find that the decision-maker wants to appoint members to the committee who are less partisan than himself, in order to facilitate the pooling of information and thereby increase the quality of decisions by the committee. This result is related to the findings of Letterie and Swank (1997) and Letterie, Swank and van Dalen (2000), who analyze the choice of a single policy advisor by a principal in order to balance the need to gain the advisor's information and the need to make the advisor's proposal credible to the final decision-maker. We also find that a balanced committee with extreme and opposite biases is acceptable to a wide range of preferences of the decision-maker, but the value of expertise from such a committee is low. By controlling the decision-making process, selective delegation provides an effective way for the decision-maker to safeguard his own interests while making use of the committee's information.

A prominent line of research on optimal delegation mechanisms in a principal-agent framework, starting from Holmstrom (1983) and Melumad and Shibano (1991), has addressed the issue of how a principal should optimally restrict the freedom of choice in delegating the decision to a single agent. By assuming that the committee's decision is

[4] Duggan and Martinelli (2001) also use a setup similar to the one presented here, with binary decisions, continuous signals and common interests among committee members. They assume voting as the decision procedure, and characterize the two-category equilibrium outcomes.

binary, we avoid this difficult issue and focus instead on how other aspects of delegation to a committee, including recruitment of committee members and control of the decision-making process in the committee, contribute to balancing the need to minimize loss of control and the need to make use of the committee expertise. Our result that a balanced committee helps strategic information aggregation and thus enhances the value of delegating to the committee, is related to the theory of advocacy by Dewatripont and Tirole (1999). They argue that it is cheaper to provide incentives for two competing advocates to gather information than for a single agent who is responsible for collecting evidence for both sides of the debate. Their argument does not hinge on eliciting private information from agents; indeed, the advantage of an adversarial system may be reduced if the collected evidence is private. More closely related is Krishna and Morgan (2001), who also deal with the issue of balance in a committee. Their model deals with two members in a committee having different preferences regarding the decision; but unlike our model, these committee members have the same data. They show that it is beneficial for a principal to consult both members (instead of consulting the one whose preference is closer to the principal) only when they are biased in opposite directions relative to the principal, because then eliciting private information is easier for the principal. In our model, balancing the committee makes the delegation more likely for a wider range of biases of the principal, because the principal can always adopt the appropriate delegation mechanism to rely on the more "loyal" member.[5]

Section 6 reports some on-going research work by Ettore Damiano and us on the role of often-observed delay in the committee decision-making process (Damiano, Li and Suen, 2008b). Using a model of repeated voting, we argue that the prospect of costly delay can bring otherwise irreconcilable committee members closer to revealing their private information in each step. The votes in each round provide endogenous information that helps committee members identify the correct decision. Even though delay does occur on the equilibrium path, ex ante welfare for each agent can be higher than what could have been achieved under a single round of voting with no delay.

[5] See also Gilligan and Krehbiel (1989) for an analysis of the open rule, modified rule, and the closed rule in legislative committees, and Chan and Suen (2007) for an application to the context of electoral competition.

Much of the existing literature on deliberations in committees assumes that such deliberations are cheap talk. Coughlan (2000) models deliberations as a non-binding straw poll prior to voting. He shows that if committee members have identical preferences, then there is full information revelation for all voting rules; and if they have preferences which are known to be different, then there cannot be full information revelation for any voting rule. Austen-Smith and Feddersen (2006) introduce preference uncertainty and constructed examples which show that cheap talk can allow full information revelation under majority voting rule but not under unanimity rule (see also Doraszelski, Gerardi and Squintani, 2003; and Meirowitz, 2007). These papers typically focus on the possibility of full-information revelation, while our work focuses on using ex post inefficient rules to help arrive at better decisions. Moreover, our model does not admit any information revelation (for sufficiently high level of conflict) in any cheap talk equilibrium, but can support an equilibrium with efficient information sharing even as the per-round delay cost goes to zero. In other words, there is a qualitative difference between modeling committee deliberations as cheap talk and modeling such deliberations as costly signals.

This article is not meant to be a comprehensive survey of the economics literature on committee decision-making, as we focus on what we believe to be the most important incentive issues using ideas from our own work. In each of the following analytical sections, sections 3 through 6, after presenting the main ideas we have a subsection that discusses some open questions and reviews the most relevant papers. We encourage readers familiar with the ideas presented in the section and more interested in the most recent developments to go directly to the corresponding subsection. The last section of the article, section 7, contains further references to recent developments in addressing other important incentive issues in committee decision-making outside our focus here, including reputation and endogenous committee formation.

2. A Model of Committee-Decision Making

Consider a committee of n members, $n \geq 2$. The committee must make a choice between a status quo and an alternative without complete knowledge about the underlying state. For convenience of presentation, we refer to the state S as being either G, which means that the alternative is "better" than the status quo, or B, which means that the alternative is "worse" than the status quo. We implicitly assume that committee members agree on the meaning of G and B. However, with information about the state dispersed in the committee, committee members may differ in their assessment of what the state is. This is discussed in full detail below.

2.1. Biases and preferences of committee members

For each $i = 1, \ldots, n$, let γ^i be the prior of member i that the state is G. Let λ_1^i be the personal cost of type I error of accepting the alternative when the underlying state is B, and λ_2^i the personal cost of type II error of rejecting the alternative when the state is G. The cost is 0 when the "right" choice is made.[6] For any belief about the state, represented by $\Pr[G|I^i]$, the probability of the state being G conditional on i's information I^i, the personal cost to i from choosing the alternative is $(1 - \Pr[G|I^i])\lambda_1^i$, while the cost from choosing the status quo is $\Pr[G|I^i]\lambda_2^i$. We refer to these payoffs as i's ex post payoffs, both because they are conditional on the updated belief that arises from new information that transpires in the committee, and because they do not include any personal cost to i that arises from investment in information gathering. Clearly, member i's personal optimal ex post decision rule is to choose the alternative if and only if

$$(1 - \Pr[G|I^i])\lambda_1^i < \Pr[G|I^i]\lambda_2^i. \tag{1}$$

We call the above i's personal optimal rule to emphasize that the decision belongs to the whole committee, and that different members will in general have different personal optimal rules.

If member i has no new information, then $\Pr[G|I^i] = \gamma^i$. In that case, member i would like to choose the alternative if and only if

$$(1 - \gamma^i)\lambda_1^i < \gamma^i \lambda_2^i. \tag{2}$$

[6] The modeling advantage of this assumption is that we can reduce preference differences to a single dimension. All that matters to the analysis is the ratio of λ_1^i/λ_2^i.

Let $k^i = (1-\gamma^i)\lambda_1^i/\gamma^i\lambda_2^i$. Then member i's optimal uninformed choice is the alternative if and only if $k^i < 1$. The parameter k^i represents the cost of false acceptance relative to false rejection to an uninformed member i. A greater k^i means that member i is more prone to rejecting the alternative. When $k^i < 1$, we say that member i is "biased toward the alternative." When $k^i > 1$, member i's optimal uninformed choice is the status quo, in which case we say that i is "biased toward the status quo." If $k^i \neq k^j$ for two members of the committee i and j, we say that the two have "conflicting interests." Conflicts between two members of the committee may arise either because they have different priors or because they have different personal costs of type I and type II errors. It turns out that in the kind of models constructed here, the source of conflicts does not matter. We assume that k^i is common knowledge in the committee.

2.2. Dispersed information in the committee

Of course, the hallmark of committee decision-making is that new information from committee members is aggregated before the joint decision is made. New information is typically dispersed in the committee. To model this, imagine that each committee member gets a signal about the state. Formally, for each state $S = G, B$, we assume that the random variable Y^i representing member i's signal is distributed on an interval $[\underline{y}^i, \overline{y}^i]$, with continuous density function $f_S^i(\cdot)$ and corresponding distribution function $F_S^i(\cdot)$. When new information in the committee is dispersed, different members have pieces of new information relevant to the decision. To capture this idea, assume that signals of different members of the committee are independent conditional on the state.[7] This is a particular way of modeling the notion that committee members have different evidence due to differences in perspectives and capabilities in evaluating the information.

We assume that the signal distributions satisfy the standard monotone likelihood ratio property (MLRP) (see, e.g., Milgrom, 1981). Let $l^i(\cdot)$ be the likelihood ratio function of Y^i, given by $f_G^i(\cdot)/f_B^i(\cdot)$. MLRP means that $l^i(\cdot)$ is strictly increasing. Intuitively, a

[7] This does not imply that the signals Y^i and Y^j for members i and j are independent because the state is unknown. In fact, conditional independence is a special case of affiliation commonly used in the auction literature. Results in the present article can be extended in a straightforward fashion to the case where signals are conditionally correlated across committee members.

higher realization of the signal Y^i is associated with a greater likelihood that the state is G instead of B.

Under MLRP, the ex post optimal decision rule is deterministic and strictly monotone for single decision-makers. For illustration, imagine that member i has access to all signal realizations $(y^1, \ldots, y^n)$ and chooses between the status quo and the alternative to minimize his expected personal cost. A "decision rule" in this context can be thought of a function p that maps any signal vector $(y^1, \ldots, y^n)$ to a probability $p(y^1, \ldots, y^n)$ of choosing the alternative. The optimal decision rule for member i minimizes the expected cost (before signals are realized)

$$(1-\gamma^i)\lambda_1^i \int p(y^1, \ldots, y^n)\Pi_{j=1}^n f_B^j(y^j)\,\mathrm{d}y^j + \gamma^i\lambda_2^i \int (1-p(y^1, \ldots, y^n))\Pi_{j=1}^n f_G^j(y^j)\,\mathrm{d}y^j. \quad (3)$$

The solution can be obtained by point-wise minimization at each signal realization, with the optimal rule given by $p(y^1, \ldots, y^n) = 1$ if

$$\Pi_{j=1}^n l^j(y^j) \geq k^i, \quad (4)$$

and $p(y^1, \ldots, y^n) = 0$ otherwise.[8] Under MLRP, each l^j is a strictly increasing function. As a result, the optimal decision rule is both deterministic, in that the probability of choosing the alternative is either 0 or 1, and strictly monotone, in that the boundary between the $p = 1$ region and the $p = 0$ region in the signal space is strictly monotone. In general, a deterministic and strictly monotone decision rule can be represented by a strictly decreasing function $M^i(y^1, \ldots, y^n)$ such that the alternative is chosen given the signal vector $(y^1, \ldots, y^n)$ if and only if $M^i(y^1, \ldots, y^n) \geq 0$. We refer to the function M^i as member's i's personal ex post optimal "decision function."

Strict monotonicity of the ex post optimal decision rule means that small changes in any member's signal realization can potentially matter to the optimal decision. Further, optimal decision functions for two members i and j with $k^i \neq k^j$ do not intersect. Indeed, if $k^i > k^j$, then $M^i(y^1, \ldots, y^n) \geq 0$ implies that $M^j(y^1, \ldots, y^n) > 0$, and conversely, $M^j(y^1, \ldots, y^n) < 0$ implies that $M^i(y^1, \ldots, y^n) < 0$. Thus, the ex post optimal choices

[8] The optimal decision rule derived here is a special case of the Neyman-Pearson lemma.

for i and j are identical for any signal $(y^1, \dots, y^n)$ such that either $M^i(y^1, \dots, y^n) \geq 0$ or $M^j(y^1, \dots, y^n) < 0$, and are different for any signal $(y^1, \dots, y^n)$ such that $M^i(y^1, \dots, y^n) < 0 \leq M^j(y^1, \dots, y^n)$.

2.3. A normal distribution example

As an example of signal structures that satisfy conditional independence and MLRP, suppose f_G^i and f_B^i $(i = 1, \dots, n)$ differ only by a location parameter. That is, $f_G^i(y^i) = g^i(y^i)$ and $f_B^i(y) = g^i(y^i - q)$ where $q > 0$. Then l^i is increasing if g^i is log-concave. Many density functions commonly used in the economics literature are log-concave. A particularly convenient example is the following.

Imagine that the underlying state is represented by a binary variable s whose value is either $q > 0$ when the state is G, or 0 when the state is B. Each member i's signal y^i is a noisy observation of the state, with

$$y^i = s + \epsilon^i, \tag{5}$$

where ϵ^i is normally distributed with zero mean and precision h^i. Then the two conditional distributions, $f_G^i(y^i)$ and $f_B^i(y^i)$, are normally distributed with common precision h^i, and means q and 0 respectively. It is straightforward to verify that the likelihood ratio l^i satisfies MLRP.

When all members of the committee have the same preferences with the ratio $k = (1-\gamma)\lambda_1/(\gamma\lambda_2)$, the ex post optimal decision rule (4) for the committee can be represented by the following threshold rule. Given observations $y^1, \dots, y^n$, with corresponding precisions $h^1, \dots, h^n$, let y denote the weighted average

$$y = \sum_{i=1}^{n} h^i y^i / h, \tag{6}$$

where $h = \sum_{i=1}^{n} h^i$. This is a sufficient statistic for the ex post decision problem. The conditional distributions of y are both normally distributed with precision h, and with mean q and 0 under states G and B respectively. Let $l(y)$ be the likelihood function associated with the conditional distributions of y. Then, condition (4) becomes simply

$$l(y) \geq k. \tag{7}$$

The above can be explicitly rewritten

$$y \geq t_*, \tag{8}$$

where

$$t_* = \frac{1}{2}q + \frac{\ln k}{qh}. \tag{9}$$

The ex post optimal decision rule for the committee is to accept the alternative if and only if the precision-weighted average y of the signals $(y^1, \ldots, y^n)$ exceeds the ex post optimal threshold t_*. Note that the optimal rule is both deterministic and strictly monotone; the decision function M that represents the boundary (in the signal space) between $p = 0$ and $p = 1$ is linear.

The ex post optimal threshold t_* includes a bias term $\ln k/(qh)$. Given the committee's common prior and preference, a greater precision h of the evidence reduces this bias term and makes the ex post optimal threshold closer to $\frac{1}{2}q$. Thus, more precise evidence allows the committee to put less weight on the bias due to prior beliefs or preferences.

Under the ex post optimal rule, the ex ante cost of type I and type II errors for the committee is

$$(1-\gamma)\lambda_1 \left(1 - \Phi\left(\sqrt{h}t_*\right)\right) + \gamma\lambda_2\Phi\left(\sqrt{h}(t_* - q)\right), \tag{10}$$

where Φ denotes the standard normal distribution function. The derivative of ex ante cost with respect to the aggregate precision h has the same sign as

$$-kt_* + l(t_*)(t_* - q). \tag{11}$$

Since the optimal threshold satisfies $l(t_*) = k$, the derivative is always negative.[9] One implication of this result is that ex ante error cost for a committee is lower the larger is the group size, provided that the optimal decision rule can be implemented.

This simple model of optimal decision-making has some other unexpected implications as well. The model shows that the optimal threshold t_* that minimizes the expected cost of the two types of errors is decreasing in k. In the context of criminal trials, efficiency

[9] If the decision rule is not chosen optimally, then an increase in the precision of the signal does not necessarily reduce ex ante error costs. See also Davis (1994).

requires the threshold of evidence needed for conviction to be lower for defendants who belong to groups with a higher prior probability of committing a crime (Farmer and Terrell, 2001). Such kind of unequal treatment before the law violates widely held notions of justice. In another application, Suen (2004) assumes that the limits of communications dictate that expert advice can only be conveyed in terms of binary recommendations, and different experts adopt different thresholds of evidence in making their recommendations. His model shows that a person who is more biased toward the status quo will optimally choose an expert who recommends the status quo more often. Such behavior, however, can lead to considerable slowdown of the learning process.

3. Conservatism in Committees

In most democratic countries, constitutional amendments must be approved by an overwhelming majority of congregations of representatives. Jury decisions in criminal trials require strong majority and even unanimity. Even if the prior of jury is biased for conviction and the jury cares little about false acquittal, the extreme conservatism of unanimity is difficult to justify. How does one explain conservatism in committees?

3.1. The free-rider problem

Committees make decision by aggregating diverse information of members. The committee decision is public because it affects the welfare of all members, but the cost of gathering information is private to the members. Greater effort by the committee as a whole produces more conclusive evidence and helps to reduce both types of error. However, a free-rider problem arises because of the public good nature of evidence. Each member in the committee disregards the benefits of more conclusive evidence to other members, leading to insufficient individual efforts in collecting information.

To apply the framework set up in the previous section, imagine a two-stage game in which committee members first simultaneously choose how much information to collect, and then the committee decision is made according to some pre-determined decision rule given all the information available to the members. To model the quantity of information collected by members, consider the normal distribution example, where each member i's signal y^i is normally distributed with precision h^i. Suppose that h^i is chosen by member i. The personal cost to member i is an increasing, convex function of h^i. Member i's total utility is the sum of the expected loss from the committee decision (from type I and type II errors) and the cost of information collection.

Suppose for now that all committee members have the same preferences. In this case, the first best outcome for the committee is unambiguously defined. It requires two conditions. First, given the members' decisions to collect information, represented by the precision levels $h^1, \ldots, h^n$, and given any resulting signals $y^1, \ldots, y^n$, the choice between the status quo and the alternative is ex post optimal. This means that the ex post optimal decision rule represented by (8) is used. Second, the decisions to collect information are ex ante optimal, given that ex post the optimal decision will be used later. It is in the ex ante stage of information gathering where the free-rider problem arises, because of the gap between the marginal value of information to the committee as a whole and the private value of information to the member that collects the information. When member i incurs a greater cost to increase the precision h^i of his signal y^i, the precision level h of the weighted average y increases. As a result, both type I and type II errors are reduced. This benefits everyone in the committee, including member i, generating the gap between the social value of information and the private value.

In the absence of a social planner that can force the members to collect the right amount of information, the first best outcome cannot be achieved. The committee can adopt the ex post optimal decision rule, but the free-rider problem implies that insufficient information will be collected by the members. In this case, even though the committee decision is optimal given collected evidence, the quality of the decision is poor. Moreover, the free-rider problem becomes more serious when there are more members in the committee if the gap between the social value of evidence and the private value increases with the size the committee.

3.2. Conservatism as a remedy

The free-rider problem in committees arises because the private benefit of gathering evidence is below its social benefit. Deviations from the ex post optimal decision rule can be a way to increase the private benefit of gathering evidence and mitigate the free-rider problem. Under the ex post optimal decision rule, the alternative is chosen if and only the weighted average y exceeds the optimal threshold t_*. Now suppose that a more "conservative" threshold is chosen: if the committee is biased toward the alternative ($k < 1$) a more conservative threshold is higher than t_*, while if the committee is biased against the alternative ($k > 1$) a more conservative threshold is lower than t_*. Li (2001) shows that a little extra caution or conservatism induces a greater private benefit of collecting evidence.

The intuition behind the claim is rather simple. The marginal value of evidence tends to be small when the committee is strongly biased either way. A conservative threshold increases the private marginal benefit of evidence by making it tougher for the committee to make the choice that is favored ex ante and thus effectively forcing a reduction in the committee's bias. To see this more clearly, suppose that a higher than the ex post optimal threshold is chosen. This reduces the chance of making the type I error of adopting the alternative when the state is B, but raises the possibility of making the type II error of rejecting the alternative when the state is G. In this case, the marginal value of information is dominated by the reduction in the type II error cost when the aggregate precision level h increases. Whether this higher threshold has a positive or negative impact on the marginal value of information depends on whether or not the committee is ex ante biased toward the alternative. When the committee is biased toward the alternative, it is more concerned with the type II error of false rejection, and thus a higher threshold increases every member's incentive to gather evidence. Therefore, a little extra caution or conservatism induces a greater private benefit of collecting evidence. Here, "caution" or "conservatism" is defined as deviation from the ex post optimal threshold against the alternative favored by the committee's prior or preference.

Although deviations from the ex post optimal threshold can induce greater participation by the committee members in gathering evidence, it is costly ex post because the threshold is suboptimal given the evidence.[10] There exists a trade-off between ex ante incentive for evidence collection and ex post optimality given evidence. For a committee of

[10] In fact, too great a deviation will eventually reduce rather than increase the private marginal value of evidence.

multiple agents, the optimal degree of conservatism is always positive. Consider increasing the threshold t slightly above the optimal level t_* when the committee is biased toward the alternative. Since t_* is ex post optimal, such a change does not affect ex post optimality at the margin, but it will provide greater incentives for evidence-collection.

The optimal degree of conservatism depends positively on the size of the committee. As the committee size increases, the gap between the social marginal value of evidence and the private value increases. A bigger committee suffers more from the free-rider problem of insufficient collection of evidence, and the benefit of a given degree of conservatism increases. Everything else equal, this will increase the optimal degree of conservatism. The free-rider problem in evidence-gathering and the conservatism remedy thus generate a scale effect on organization behavior. This scale effect explains the often-made observation that larger organizations tend to have a greater organization inertia, such as slowness in discarding old inefficient organization habits.

The result about the benefit of conservatism has implications for jury trials. Although in the adversarial system evidence is competitively provided by the defendant and the plaintiff, jurors must spend effort in digesting arguments put forth by the lawyers on both sides and following instructions by the trial judge in deliberating the verdict. To the extent that such effort costs individual jurors but benefits the whole jury, a free-rider problem exists in the jury fact-finding process. If for serious crimes the prior of the jury favors conviction—perhaps a reflection of the jury's belief in law enforcement or the high conviction rate in actual juries for such crimes, then a higher conviction threshold than what is ex post optimal helps alleviate the free-rider problem in the jury. An implication is that the conviction threshold used by juries in criminal cases should be higher than those used by juries in civil lawsuits, because juries in criminal cases are typically larger in size and therefore suffer more from the free-rider problem. This is consistent with the fact that the threshold of "guilt beyond reasonable doubt" used in criminal cases is higher than the threshold of "preponderance of evidence" in civil lawsuits. Another implication is that the conviction threshold used by juries in the adversarial system of Britain and America should be higher than in similar trials in the inquisitorial system of continental Europe, where a single judge decides and the free-rider problem is therefore absent.

The basic result that a little extra caution helps mitigate the free-rider problem in committees extends to the case of committee members having different preferences. Consider an extreme case of heterogeneity: a committee has two members, 1 and 2, with member 1 concerned only with type I error and member 2 concerned only with type II error. When the committee is dominated by member 2's concerns for type II error, a higher than the ex post optimal threshold can yield greater total incentives in gathering evidence. Although member 1 reduces the information-gathering effort as the result of a higher threshold, member 2's increase in information-gathering effort more than compensates for the decrease in member 1's effort.

Conservatism is a commitment by a group to induce more effort in gathering information in the presence of the free-rider problem and in the absence of centralized effort in gathering evidence. The informational requirement for this commitment is not stringent, because implementing a conservative rule requires neither observation of the precision level of the summary evidence nor any information about the individual levels of precision or effort. If the group could commit to decision rules that depend on aggregate level or even individual levels of precision and effort, then the first best outcome may be achieved through a threat. Consider the following decision rule: if the individual level of precision is at least h_*, the first best individual level of precision, then choose the decision according to the ex post optimal standard; otherwise, choose the decision that maximizes the expected loss given the evidence. However, such schemes are seldom observed in situations involving collective decisions. The reason may be that their implementation requires the group to monitor the precision level or the effort level, and this is too strong to be realistic.

3.3. Discussion

Persico (2004) compares different decision rules in a private-information model that combining information aggregation and information acquisition. In his model, each agent must incur a fixed cost to learn a binary signal, and a mechanism designer can choose a subgroup to make the decision. He shows that decision rules close to veto power of individual members in the subgroup are optimal only if the signals are sufficiently accurate. Moreover, in the class of threshold voting rules that select the alternative over the status quo only if the number of votes for the alternative exceeds a certain minimum, the optimal rule coincides with the ex post efficient threshold rule. However, Gerardi and Yariv (2008) show that the conclusion changes once we allow for a broader class of voting rules. Using a mechanism design approach, they show that the optimal mechanism is in general not ex post efficient in order that agents have greater incentive to acquire information.

The free-rider problem in small committees is related to the rational voter ignorance problem pointed out by Downs (1957). In large elections, each voter has a negligible chance of affecting the electoral outcome. Therefore, each voter has little incentive to become well-informed about the issues to be decided. When the cost of information acquisition is small, it is not clear whether aggregating diverse information from a large number of poorly informed voters will produce desirable outcomes. A recent literature deals with the question of whether the Condorcet jury theorem continues to hold when acquiring information is costly. Mukhopadhaya (2005) shows that in a symmetric mixed strategy equilibrium, as the number of committee members increases, each member chooses to collect information with a smaller probability. He finds examples in which, using the majority voting rule, a larger committee makes the correct decision with a lower probability than does a smaller one. Koriyama and Szentes (2007) consider a model in which agents choose whether or not to acquire information in the first stage, and then the decision is made according to an ex post efficient rule in the second stage. They show that there is a maximum group size n^P such that each member collects evidence with probability 1 if $n \leq n^P$. Moreover, the optimal group size is either n^P or $n^P + 1$. Thus the Condorcet jury theorem fails beyond the optimal group size. However, they also show that the efficiency loss from an excessively large group is relatively small. In a related framework but with the quality of information as a continuous choice variable, Martinelli (2006) shows that if the marginal cost of information is near zero for nearly irrelevant information, then there will be effective information aggregation despite the fact that each individual voter will choose to be very poorly informed.[11] Martinelli (2007) assumes that voters have heterogeneous costs of information acquisition. He finds that if the support of the

[11] Whether such an outcome is efficient after taking into account the cost of information acquisition is shown to depend on the second derivative of the cost function for nearly irrelevant information.

cost distribution is not bounded away from zero, then there is an equilibrium with some information acquisition even for arbitrarily large elections. Moreover such an equilibrium dominates an equilibrium with no information acquisition in terms of aggregate welfare. These results suggest that large elections can be effective, albeit not necessarily efficient, means of aggregating information despite the free-rider problem.

Gathering information about a new alternative choice in a committee often takes time. This adds a dynamic dimension to the free rider problem described above, and creates a new countervailing "encouragement effect," as the prospect of free-riding on future information-gathering efforts by other members encourages each individual member to continue to gather information beyond the point where he would have stopped if he were alone.[12] It is straightforward to write down a dynamic version of the information acquisition model in section 3.1, where individuals decide how much additional information to acquire in each period under some commitment that makes the decision once the learning has stopped. Whether there is too much or too little experimentation under the ex post optimal decision rule, and whether modifications of the decision rule help provide the right incentives are potentially important questions worth further investigation.

4. Strategic Information Aggregation

In this section, we take the information available to each committee member as exogenously given. Unlike in the previous section, we assume that the information is private. A host of new issues arises when committee members have conflicting interests because they will want to misrepresent their private information to influence the committee decision. Most of these issues can be adequately illustrated by having two committee members.

4.1. Voting as equilibrium garbling

Since information is private, committee decisions are made on the basis of members' reports of their private data. Let us first consider a Bayesian game where the two members report

[12] The encouragement effect is first analyzed in Bolton and Harris (1999), in a strategic experimentation framework with private adoption decisions and public learning. See also Bergemann and Valimaki (1992) and Keller, Rady and Cripps (2005).

r^1 and r^2 simultaneously after learning their private evidence y^1 and y^2, and the decision is made according to a decision rule of "accept the alternative if and only if $M(r^1, r^2) \geq 0$," where M is a strictly decreasing decision function. It is easy to see that for any strictly decreasing M, truthful reporting is not an equilibrium strategy as long as $k^1 \neq k^2$. This is because $k^1 \neq k^2$ implies that at least one of the two personal optimal decision functions M^1 and M^2 differs from the committee decision function M. Suppose member 2 always reports his observation y^2 truthfully. If member 1 submits report r^1, the alternative is chosen if $M(r^1, y^2) \geq 0$. But conditional on y^1, member 1 prefers to the alternative whenever y^2 satisfies $M^1(y^1, y^2) \geq 0$. When M differs from M^1, reporting $r^1 = y^1$ is not optimal for member 1.

In Li, Rosen and Suen (2001) we show that, given any deterministic and strictly monotone committee decision rule, there exists no manipulation equilibrium where members use invertible reporting strategies that allow perfect inference of their private data. This result illustrates the incentives to garble private information in committee decision-making. Since our argument depends only on the local characteristics of the reporting strategies, there exists no equilibrium with partially invertible strategies (i.e., reporting functions that are invertible for some interval in the support of the evidence) either. Garbling occurs almost everywhere. Since any Pareto-optimal ex post decision rule is necessarily deterministic and strictly monotone, the garbling result implies that there is no mechanism that can implement a Pareto-optimal ex post decision rule.

The above observations can be made precise and strengthened further from a mechanism design perspective. Formally, a "decision mechanism" here is a commitment by member 1 and member 2 to a "report space" for each member that defines all the reports he can choose, and a "committee rule" that maps a vector of reports to a decision. The Bayesian reporting game considered above defines a "direct decision mechanism," because each member's report is confined to the range of the observations. One can easily imagine "indirect mechanisms" where members' reports are not restricted to this range. For example, a voting procedure is an indirect mechanism because the report space for each committee member consists of two votes, yes and no. Whether direct or indirect, if a decision-making mechanism has an equilibrium, then the equilibrium defines an "outcome," a mapping from the data space to the decision that is a combination of the equilibrium strategies and the committee rule. A deterministic and strictly monotone outcome has a boundary between the acceptance and rejection regions that is a strictly decreasing function in the signal space. Another possibility is a deterministic but everywhere weakly monotone outcome, represented by a boundary that is a decreasing step function. This is called a "partition outcome," because continuous data of each member are transformed into ordered ranks or categories through partitioning by thresholds. Categorizing data is a particular form of information garbling that coarsens information in a natural way and prevents full revelation of private evidence. Applying the revelation principle (Myerson, 1979), we can exclude any deterministic outcome whose boundary has a strictly decreasing segment as an equilibrium outcome.

Thus, committee decision-making can be accomplished only by categorizing private data. Coarsening information through ordered categories controls private incentives to exaggerate the strength of one's private evidence while at the same time affording opportunities to pool everyone's evidence. There is a trade-off between information sharing and manipulation. How this trade-off is resolved depends on the degree of a priori conflict among the members. The greater the latent consensus among members, the greater are the opportunities for presenting private data in finer categories and greater detail.

The result of equilibrium garbling gives a strong sense that we derive "voting" with categories as a necessary method to achieve consensus in committee decision-making. This is the clearest when partition outcomes involve two categories. In this case, each member has two votes, "accept" or "reject" (the alternative), and a rule stipulates how many votes are needed to accept the alternative. More generally, voting procedures can allow each member to choose an integer score from a given scale (say, 1 to 10), and a pre-specified scoring rule aggregates the scores and compares the sum to a given threshold (say, 9.5) to yield a committee decision. The equilibrium outcomes in these generalized voting games correspond to partition outcomes with multiple categories. Economists often complain that voting procedures are imperfect means of collective decision making because they do not allow individuals to express their intensity of preferences. Our result shows that it is precisely the inability to fine-tune the reports about members' preferences that allows meaningful information sharing to take place in the committee.

4.2. Information manipulation and information sharing in voting games

Let us now consider voting games in which each member chooses between "accept" or "reject." Although restrictive, these games nicely illustrate the main strategic issues that arise from information manipulation and information sharing in a committee. With two votes for each of the two committee members, two different voting procedures are possible: "unilateral acceptance" (where the alternative is accepted if there is at least one "accept" vote) and "unilateral rejection" (where acceptance requires two "accept" votes). Each procedure defines a different voting game, with an equilibrium in which each member votes "accept" when his observation is above a threshold. An equilibrium with such threshold strategies results in a partition outcome with two categories.

To see how to characterize an equilibrium in a two-category voting game, consider unilateral acceptance; the case of unilateral rejection is similar.[13] An equilibrium is a pair of thresholds t^1_* and t^2_* that satisfy the following two indifference conditions. First, given that member 2 adopts the strategy of voting "accept" if and only if $y^2 \geq t^2_*$, member 1 is indifferent between the two votes when his own signal is $y^1 = t^1_*$. The second indifference condition is symmetric. Each condition may be thought of as a reaction function. For example, the first condition defines the optimal acceptance threshold of member 1 for any acceptance threshold of member 2. Under MLRP, the two reaction functions are downward sloping.[14]

A few comparative statics results for the voting game illustrate the tension between information manipulation and information sharing. First, if Y^1 and Y^2 have the same conditional distributions, then $k^1 < k^2$ implies $t^1_* < t^2_*$. For the same observation value $y^1 = y^2 = y$, member 1 votes for acceptance while member 2 votes for rejection if y is between t^1_* and t^2_*. Therefore $|t^1_* - t^2_*|$ can be thought of as the "area of disagreement" between the two members. Second, the area of disagreement increases as the conflict of interests, $|k^1 - k^2|$, increases. As member 2 becomes more biased toward rejection and his standard

[13] The equilibrium conditions can also be established using the pivotal voter argument.

[14] Equilibrium exists under appropriate boundary conditions on the likelihood ratios. A strengthening of MLRP implies uniqueness of the equilibrium. This condition is satisfied if the conditional distribution functions of Y^1 and Y^2 are normal.

for acceptance increases, member 1 counters by lowering his own standard. This induces member 2 to raise t^2_* further. The increase in t^2_* can be decomposed into two parts: one due to a shift of 2's reaction function, and the other due to a move along 2's reaction function because t^1_* decreases. The second part shows that the area of disagreement in committee decision-making is larger than that implied by inherent conflicts in preferences: strategic manipulation by one member leads to counter-manipulation by the other. Conflicts lead to the exaggeration. Third, although conflicts cause manipulation, incentives to exaggerate evidence are balanced in equilibrium by incentives to share information. If a member makes the decision alone, the optimal decision is acceptance if and only if his evidence y^i $(i = 1, 2)$ exceeds a threshold $\hat{t}^i$ determined by $l^i(\hat{t}^i) = k^i$. This threshold $\hat{t}^i$ is lower than t^i_*: when i observes evidence y^i between $\hat{t}^i$ and t^i_*, he votes for rejection in the committee even though he would have chosen acceptance if he were the only decision-maker. Member i thus utilizes the information of the other member by casting the decisive "accept" vote less frequently. Note that this is true independently of member i's preferences. Even if member i is strongly biased toward acceptance, the need to utilize the other member's information still makes him more "conservative" towards acceptance. Moreover, Li, Rosen and Suen (2001) establishes that the more conservative member always benefits more from committee decision-making with a unilateral acceptance rule than from decision-making based on his own private information. Indeed, as long as the degree of conflict is not too large, each member prefers making the decision through the committee to dictating the outcome on the basis of his own information.

Incentives to share information under conflicting interests can also be examined by considering how voting behavior changes when quality of the observation received by one member, say 1, becomes higher. If evidence were public, higher quality data receive greater weight in the decision rule. But since evidence is private, changes in weights can be easily undone by information manipulation. Instead, changes in information quality changes equilibrium thresholds. An increase in the quality of member 1's signal causes t^1_* to decrease and t^2_* to increase. The interpretation is straightforward. Voting for acceptance decides the verdict regardless of the value of the other member's observation. Voting for rejection, on the other hand, defers the decision to the other member. When member 1 gains access to data of a higher quality, member 2 takes advantage of the improved information by raising t_*^2 and deferring the decision to member 1, so member 1 is decisive more often.

4.3. Discussion

Comparing the equilibrium with unilateral acceptance to the equilibrium with unilateral rejection in the two-category voting games, one may conclude that requiring two votes for acceptance instead of one is a more "stringent" standard of proof. But this statement ignores strategic responses to the voting procedure. When unanimity in acceptance is required, each member lowers his threshold for acceptance and votes for acceptance less cautiously, because he knows that the other member may have information that will lead to a vote against acceptance. On the other hand, if acceptance is unilateral, each member is more cautious in casting a vote for acceptance, because such a vote would be decisive regardless of the other member's information. The comparison between unilateral acceptance and unilateral rejection depends on the precise shapes of the conditional distributions of Y^1 and Y^2. In particular, unanimous acceptance does not necessarily lead to lower acceptance rates. This comparison of voting procedures complements the works of Sah and Stiglitz (1986; 1988), who consider committees without the strategic manipulations that arise from conflicting interests. It is also consistent with the existing literature showing why the Condorcet jury theorem fails when strategic considerations play a role in voting (Austin-Smith and Banks, 1996; Feddersen and Pesendorfer, 1998). In particular, Feddersen and Pesendorfer (1998) show that a unanimous conviction rule in jury decisions may lead to higher probability of false conviction as well as false acquittal than a simple majority rule, and the probability of convicting an innocent defendant may increase with the size of the jury.

In this section we have shown that garbling must occur in equilibrium under any decision rule as long as there are conflicts in the committee, and we have established a few characterization results for deterministic and weakly monotone outcomes (partition outcomes). It can be shown that within the class of deterministic and weakly monotone outcomes, the Pareto optimal outcome allows the maximum number of partitions in each member's signal realization, which is in turn determined by the extent of conflicts in the committee. However, we do not know if stochastic outcomes (i.e., probabilistic choice between the status quo and the alternative) or non-monotone outcomes are Pareto dominated by partition outcomes. More importantly, we have not asked the question of what is the optimal mechanism that maximizes a weighted sum of ex ante utilities of committee members. This question is difficult to answer, because unlike the standard mechanism design problem, it involves common values and no side transfers. Indeed, even with one-sided private information as in Crawford and Sobel (1982), the question of optimal mechanisms is already theoretically challenging. Recently, Alonso and Matouschek (2008) have made progress on this problem, by restricting the preferences of the sender (agent) and the receiver (principal).[15] Their characterization of the optimal delegation mechanism holds some promise for answering the question of optimal mechanisms in committees.

5. Management of Committees

In non-market organizations, expert opinions on choices facing a principal decision-maker are frequently provided by standing committees that are recruited and put in place before the choices have to be made by the principal. The discussion so far indicates that committees are effective, albeit imperfect, mechanisms for aggregating diverse information. In this section we analyze how a principal decision-maker can use delegation and recruitment strategies to control information manipulation and safeguard his own interests, when conflict of interests exists both between the principal and the committee members and among the committee members themselves.

5.1. Delegating decisions to a committee

When decisions are delegated to a single agent, this agent decides according to his own preferences, which do not necessarily accord with those of the principal. One way to control manipulation is to form a committee and play one committee member against

[15] Kovac and Mylovanov (2006) provide sufficient conditions for the optimal mechanism to be deterministic with quadratic preferences. They also construct examples that show in general optimal mechanisms involve randomization even though the receiver's decision is a continuous choice.

another. In other words, conflict of interests among the committee members themselves may prevent the principal from being dominated by any single committee member. An additional advantage from using more than one agent arises from potential gains from pooling diverse information. We illustrate these issues in the simplest terms by considering a committee of two members.

An uninformed principal D considers whether to dictate the outcome or to delegate to a committee of two—member 1 and member 2. If delegation occurs, member 1 and member 2 face a problem of strategic information aggregation. Such a game is analyzed in section 4. For the present purposes, we focus on simple two-way voting equilibria, corresponding to the type of two-way voting procedures. In "unilateral acceptance," the alternative is accepted if there is at least one "accept" vote. In "unilateral rejection," acceptance requires two "accept" votes. These two voting procedures have natural interpretations as two different "delegation mechanisms."

The two delegation mechanisms define two different voting games between member 1 and member 2. As described in the previous section, the equilibrium associated with each mechanism is characterized by a pair of acceptance thresholds, one for each member. In the voting game with unilateral acceptance, each member $i = 1, 2$ votes "accept" if and only if $y^i \geq t^i_*$; in the voting game with unilateral rejection, each member i "accept" if and only if $y^i \geq t^i_{**}$. Clearly, $t^i_{**} < t^i_*$ for each member $i = 1, 2$.

Given the equilibria, the principal's choice of whether or not to delegate the decision to the committee depends on his own preference. We establish the following result: principal D prefers an uninformed acceptance decision to delegation under unilateral acceptance if he is sufficiently biased toward the alternative (if his preference ratio k^D is sufficiently small), prefers an uninformed rejection decision to delegation with unilateral rejection if he is sufficiently biased against the alternative (if k^D is large enough), and prefers delegation if k^D is in between these critical values.[16] Thus, the principal optimally chooses

[16] The range of delegation should be viewed as a conservative bound. With arbitrary distributions satisfying the monotone likelihood ratio property, we cannot rule out the possibility that principal D prefers delegation under unilateral rejection to making an uninformed acceptance decision when he is sufficiently biased against the alternative. In other words, unilateral rejection may be preferred to unilateral acceptance when k^D is low.

not to delegate if his preference biases are too large in either direction. In contrast, an unbiased principal (with $k^D = 1$) will always prefer delegation to making an uninformed decision. Further, any principal with preferences k^D between k^1 and k^2 will always prefer delegating the decision to the committee of member 1 and member 2 to making the decision himself. In other words, delegating the decision to the committee dominates making an uninformed decision as long as the two committee members have opposite biases relative to the principal.

Since a committee of two members has access to more information than a single agent, the committee has a greater "influence" than either of its member by himself. For a principal who is willing to delegate the decision to member 1, adding any member 2 cannot reduce his willingness to delegate the decision. In the worst cases when member 2 has extreme preferences (k^2 is either close to zero or infinity), this extreme member can be rendered irrelevant by an appropriate choice of the delegation mechanism for the committee. For example, if member 2 is strongly biased toward the alternative (k^2 close to zero), the principal can choose unilateral rejection so that delegating the decision to the committee of member 1 and member 2 is equivalent to delegating to just member 1.

A wide use of delegation by different types of principals is ensured if the two committee members have extreme and opposite preferences. Such a balanced committee of radicals will have a wider delegation range than a one-sided committee of moderates. To see this, consider a committee of two members with moderate preferences and $k^1 < k^2$. Any principal with preference k^D sufficiently biased toward rejection prefers making an uninformed decision to delegating the decision to this committee. This problem can be overcome by recruiting a member 2 with a higher k^2. However, doing so without changing k^1 will make the committee less acceptable to principals who are biased toward acceptance. The solution is to reduce k^1 while increasing k^2, that is, recruit members who are more extreme in opposite directions. A balanced committee has a greater use of delegation than a one-sided committee because the principal can specify the voting procedure (unilateral acceptance or unilateral rejection) in delegating to the committee. When the principal is biased toward the alternative, his natural ally is the member who has a similar bias, and the delegation mechanism of unilateral acceptance protects the interests of the principal by giving the control of decision-making in the committee to his ally. Similarly, when the principal is biased toward the status quo, he relies more on the member who is relatively biased toward the status quo, and the procedure of unilateral rejection protects the interests of the principal.

However using a balanced committee with extreme preferences is not costless. Such a committee may be very influential in the sense that it is acceptable to a wide range of principals, but the value of its expertise, while positive, is generally low. Greater conflict of interests within a balanced committee of extremists makes information transmission less efficient. As a result the quality of the committee's decision deteriorates. Thus, there is a trade-off between influence and value for a committee. A committee that is acceptable to a wide class of principals may not produce decisions that are highly valuable.

5.2. Recruiting committee members

For a given principal D, the ideal committee should have both members with preferences identical to D's. However, the analysis above shows that having balanced preferences in the committee makes it widely acceptable to different principals. Thus, there is a benefit of having some degree of conflict in the committee. Consider, then, a situation in which principal D already has member 1 at hand and is thinking about recruiting a member 2 to form a committee with member 1. Suppose that $k^D < k^1$ so that member 1 is relatively biased against the alternative.

Does principal D want recruit a second member with the same preference as himself ($k^2 = k^D$)? The answer is to this question is no. The reason is that having $k^2 = k^D$ will force the existing member 1 to be more extreme and therefore distort member 1's information. Conflict of interests exists between the two members within the committee, as well as between member 2 and the principal. By recruiting a second member with the appearance of neutrality, the principal can reduce member 1's acceptance threshold and make member 1 less manipulative with his information. Indeed, with some additional mild assumptions it can be shown that the best member 2 for the principal to recruit is one with k^2 between k^D and k^1, regardless of whether the delegation mechanism is unilateral acceptance or unilateral rejection.

A related recruitment issue is whether it is possible that the conflicts between D and member 1 are so large that D wants to recruit a member 2 and delegate to member 2 alone. If the delegation mechanism cannot be chosen, then having a committee of two members can be worse than having just one for the principal. For example, a principal strongly biased toward the alternative prefers delegating to a single agent with moderate preferences than to a committee consisting of this agent and another one with an extreme bias against the alternative, if the committee has to use the unilateral rejection mechanism. This is because the unilateral rejection mechanism allows the second member of the committee to make the rejection decision by himself with little input from the moderate member. However, this problem can be avoided if the principal instead chooses the unilateral acceptance mechanism. In the worst scenario, the committee member with the extreme bias against the alternative would not affect the decision at all, and in general the member's information will be used in the decision. It can be established that as long as the principal can select the voting procedure, a committee of two is preferred to the single moderate member.

The result that the principal wants to recruit a new committee member with the appearance of neutrality does not depend on whether the principal is informed or not. Even if the principal is equally well informed as the member, he has an incentive to delegate the member to sit in the committee on his behalf in order to reduce the conflict of interests within the committee. By appointing someone with preferences more similar to other committee members' preferences, the principal commits himself to be less partisan and thus improves the pooling of information within the committee. While we have illustrated this point under the assumption that the two members play a simple two-way voting game, the same conclusion holds when more sophisticated mechanisms are considered. For example, a principal or his delegate may make a recommendation and a higher authority ultimately determines the outcome based his own private information and on the recommendation. As long as conflict of interests exists between the principal and the higher authority, our theory suggests that the principal will recruit a delegate whose preferences are closer to those of the higher authority. This ensures that the higher authority will be less skeptical of the recommendation, and therefore rely less on his own preferences and more on the recommendation in making the final decision. The principal benefits from the resulting improvement in the efficiency of information aggregation.

5.3. Discussion

While a committee of members with opposite but extreme biases is acceptable to a wide range of principals, the value of advice from such committee is low because of information manipulation. The dysfunctional aspect of "competition of idea" has received an interesting treatment in a recent paper by Dessein (2007). By introducing costly state verification, Dessein (2007) formalizes the notion that committee decision-making can be costly as self-interested members make cheap-talk proposals that require detailed examination and discussion. Under some circumstances, having an authoritarian leader to make the decision can reduce rent seeking by people with bad ideas, while preserving the incentive of those with genuinely good ideas to submit their proposals.

The delegation mechanisms considered in this section require the principal to specify the game that the committee members play upon delegation. At the same time, in order for the principal not to have incentives to overrule the decision reached by the committee, it is important that the principal does not monitor how the decision is reached. It is interesting to study how principals can balance the need to minimize loss of control and the need to make use of expertise in the committee in the absence of the commitment to no monitoring and the commitment of not overruling. Future research along this line may benefit from a few recent papers in the literature on optimal delegation with a single agent that have tried to address the issue of commitment (Dessein, 2002; Marino, 2007; Mylovanov, 2008). They center on the idea that the principal may not be able to commit to not vetoing the agent's choice when there is a default action that the principal could take ex post. Mylovanov (2008) reconciles contradictory findings in Dessein (2002) and Marino (2007) by showing that if the principal can choose a default action ex ante, then the optimal delegation mechanism under full commitment (for the case of uniformly distributed signal of the agent and quadratic preferences for the agent and the principal) can be implemented veto-free, so that the inability to commit to not taking the default action ex post is inconsequential. Li and Li (2009) take a different approach to the issue of commitment: in their model, the principal can commit to not changing the agent's choice to an action outside the set of actions delegated to the agent, but is free to switch it to another action in the set. They show that this form of limited commitment is binding, in that it yields a different solution from the one that is optimal under full commitment. Furthermore, in an optimal mechanism under limited commitment, the principal specifies fewer actions delegated to the agent in order to relax the incentive compatibility constraints due to limited commitment.

Much of the literature on delegation and the management of committees assumes that ultimate authority lies in a principal decision-maker. In many situations, however, committees have the ultimate decision-making power, and the question of "management of committees" is really about how interested parties can influence their decisions. Caillaud and Tirole (2007) take a first step in addressing this question. They study how a proposer builds consensus for his proposal by using a persuasion cascade to distil the information to the key members of the committee. More work in this area remains to be done.

6. Deliberations and Delay

A common complaint about committee decision-making is the lengthy delay it often entails. When individuals with private information and diverse preferences have to make a collective decision, disagreement within a committee is almost inevitable. In principle, disagreement driven by private information can be satisfactorily resolved if there is an efficient mechanism for sharing information. Disagreement driven by preference differences, however, has no simple solution. Moreover, individuals have incentive to disguise preference-driven disagreement as information-driven disagreement, making efficient information aggregation more difficult. Just as an ex-post inefficient decision rule can be an effective remedy to the free-rider problem, we argue that seemingly inefficient delay in the decision-making process can be an effective remedy to the information aggregation problem. Costly delay makes people more forthcoming about their private information, and repeated deliberations produce endogenous information that allow better decisions to be made. We illustrate these points using the basic model discussed in this article as well as a model of repeated voting developed from on-going research by Damiano, Li and Suen (2008b).

6.1. Repeated information aggregation

In section 3 we see that voting is a method to control the incentives to manipulate information when committee members have conflicting interests. An individual who casts one round of "accept" or "reject" vote can be interpreted as partitioning his continuous information into two distinct categories. When the degree of conflict is not too strong, equilibrium can generally allow a finer partitioning of information that will produce a more efficient use of information. One way to achieving such kind of finer partitioning of information is through repeated voting.

Consider, for example, a voting game with the following rules. Each of the two committee member can vote to accept or reject an alternative in the first round. If the votes agree, then that decision is implemented immediately. If the votes disagree, then they move to the second and final voting round. In this second round, the decision is made by unilateral acceptance, say. Assume for the moment that there is no cost of delay.

Depending on the parameter values, this repeated voting game can support a three-category or four-category partition equilibrium. For example, suppose $k^1 < k^2$, and $k^2 - k^1$ is not too small. Then, equilibrium in the two-round voting game takes the following form. In the first round, each committee member votes to accept if and only if $y^i \geq t^i$. Since member 1 is more willing to accept the alternative than member 2, we say that there is a "reverse disagreement" when the first member votes to reject and the second member votes to accept. Upon a reverse disagreement, member 1 learns that $y^2 \geq t^2$, so he becomes even more prone to accept. On the other hand, member 2 learns that $y^1 < t^1$, so he becomes even more prone to reject. In other words, a reverse disagreement widens their initial conflict of interests. If the $k^2 - k^1$ is not too small, the subgame following a reverse disagreement may allow no information aggregation at all; that is, member 1 always votes to accept and member 2 always votes to reject following a reverse disagreement. Another type of disagreement is a "regular disagreement," in which member 1 votes "accept" and member 2 votes "reject" in the first round. A regular disagreement tends to moderate the degree of conflict between the committee members. So, if member 1 is initially indifferent between "accept" and "reject" (i.e., $y^1 = t^1$), he will strictly prefer to reject following a regular disagreement. In other words, there is a range of signal values $y^i \in [t^1, s^1]$ such that member 1 would switch from voting "accept" in the first round to voting "reject" in the second round following a regular disagreement. Similarly, there is a range of signal values $y^2 \in [s^2, t^2]$ such that member 2 would switch from "reject" to "accept" following a regular disagreement. Thus the way information is aggregated in this game depends on whether y^1 belongs to $[\underline{y^1}, t^1)$, $[t^1, s^1)$ or $[s^1, \overline{y^1}]$ and on whether y^2 belongs to $[\underline{y^2}, s^2)$, $[s^2, t^2)$ or $[t^2, \overline{y^2}]$.[17] Although there is no exogenous information arriving between the two voting rounds, the endogenous information from observing how committee members vote in the prior round leads to more efficient information aggregation than two-way voting.

6.2. Using costly delay to mitigate information distortion

When delay is costless, the value of the thresholds t^i and s^i in the model of repeated voting described above will be identical to those of a three-category equilibrium as in Li, Rosen and Suen (2001). In that case, repeated voting is only one among many possible mechanisms to implement the three-category equilibrium.[18] The situation changes when delay is costly. Even though delay is wasteful ex post, the prospect of costly delay prompts committee members to be more forthcoming about their private information.

In our previous example, the probability that the decision has to be delayed until the second round is equal to $F^1_S(t^1)[1 - F^2_S(t^2)] + [1 - F^1_S(t^1)]F^2_S(t^2)$ in state $S = G, B$. When delay is costly, committee members have incentives to lower the probability of delay by changing their thresholds. Suppose $k^1 < k^2$. A sufficient condition that member 1 can reduce the probability of delay by raising his threshold t^1 is $F^2_G(t^2) > \frac{1}{2}$. Similarly, a sufficient condition that member 2 can reduce delay by lowering his threshold t^2 is $F^1_B(t^1) < \frac{1}{2}$.[19] Recall that strategic information manipulation leads to exaggeration: the member who is predisposed to acceptance chooses an acceptance threshold that is too low

[17] When $k^2 - k^1$ is sufficiently small, some information aggregation is still possible even following a reverse disagreement. In that case, the outcome is the same as a four-category information-partition equilibrium.

[18] Li, Rosen and Suen (2001) discuss other possible mechanisms in a static framework.

[19] More precisely, starting at an equilibrium without delay cost, a marginal increase in delay cost raises equilibrium t^1 if and only if

$$\frac{1 - 2F^2_G(t^2)}{2F^2_B(t^2) - 1} < \frac{(F^2_G(t^2) - F^2_G(s^2))\lambda^2_2}{(F^2_B(t^2) - F^2_B(s^2))\lambda^2_1};$$

for efficient information aggregation, while the member who is predisposed to rejection chooses an acceptance threshold that is too low. If $F_G^2(t^2) > \frac{1}{2} > F_B^1(t^1)$ when delay cost is zero, a marginal increase in delay cost will raise t^1 and lower t^2. Moreover, because choice of thresholds is a supermodular game, it is also the case that equilibrium s^1 will rise and equilibrium s^2 will fall. In other words, member 1 raises each of his acceptance thresholds while member 2 lowers his. Strategic information manipulation by committee members is lessened as they try to avoid the delay costs entailed by evidence exaggeration.

Li, Rosen and Suen (2001) show that expected error cost in the three-category equilibrium falls for each committee member as they simultaneously reduce information manipulation by moving their acceptance thresholds toward the direction of each other. Thus ex post inefficient delay can be viewed as one way to improve the ex ante quality of decisions reached by a committee. To be sure, the improvement in quality of decisions has to be balanced against the direct costs of delay. A general cost-benefit analysis is cumbersome as each additional round of voting increases the number of thresholds to be considered. Nevertheless our recent work (Damiano, Li and Suen, 2008b) establishes that, in a simplified setting with binary signals, repeated voting with small delay costs will unambiguously raise the ex ante welfare of each committee member when the degree of conflict is moderate.

Damiano, Li and Suen (2008b) consider a model with three states. If the state is A, both member 1 and member 2 prefer to accept the alternative. If the state is R, both prefer to reject. The third state C is a conflict state, in which member 1 prefers to accept while member 2 prefers to reject. Member 1 can distinguish state A from states $\{C, R\}$, while member 2 can distinguish state R from states $\{C, A\}$. Thus, member 1 who knows that the state is A knows that accepting the alternative is the correct decision, and he would like to convince member 2 that the state is indeed A. However, if the probability of the conflict state is sufficiently high (greater than $\frac{1}{3}$, in particular), member 1 who is uninformed and who knows only that the state is either C or R would also like to mislead

and it lowers equilibrium t^2 if and only if

$$\frac{1 - 2F_G^1(t^1)}{2F_B^1(t^1) - 1} > \frac{F_G^1(t^1)\lambda_2^1}{F_B^1(t^1)\lambda_1^1}.$$

member 2 into believing that the state is A. We show that in this situation, any incentive compatible mechanism without side transfers cannot dominate making the decision by a simple coin toss. No information aggregation takes place even though sometimes it is in the interest of the committee members to agree if they could share their private information.

Introducing repeated voting with costly delay dramatically alters the negative conclusion. In a game in which voting continues indefinitely as long as there is disagreement between the committee members, there is an equilibrium in which an informed committee member always votes for the correct decision, while an uninformed committee member randomizes between voting to accept and voting to reject. In this equilibrium, a reverse disagreement (member 1 votes to reject while member 2 votes to accept) leads players to conclude that the state is the conflict state, and the ensuing sub-game is similar to a war of attrition. A regular disagreement (member 1 votes to accept while member 2 votes to reject) leads the uninformed players to lower their belief that the state is the conflict state. Hence, if member 1 does not know whether the state is C or R, he votes to accept with a decreasing probability in the next round following a regular disagreement. It is possible to show that if the state is either A or R, the Pareto-preferred decision is always reached in a finite number of rounds. Thus the possibility of repeated voting with delay costs allows information to be aggregated efficiently in an environment where information aggregation is otherwise impossible. We also show that when delay cost per voting round is sufficiently low and when the degree of conflict is not too large, the ex ante welfare of each committee member is strictly higher than that is achievable under any general mechanism that does not allow for transfers or delay costs.

6.3. Discussion

It is well-known that in the one-sided private information framework of strategic information transmission such as Crawford and Sobel (1982), unrestricted, costless back-and-forth communication, or long cheap talk, can improve upon the equilibrium outcomes of a single-round of communication. A general treatment of the case of discrete actions and type spaces can be found in Aumann and Hart (2003). For the special case of uniformly distributed private signal and quadratic preferences of the cheap talk model of Crawford and Sobel (1982), various authors have shown that multiple rounds of communication can help (Krishna and Morgan, 2004; Blume and Board, 2007; Goltsman, Horner, Pavlov and Squintani, forthcoming). As evidenced in these papers, a full analysis of the general case is exceedingly difficult. One expects that the analytical difficulties carry through to our framework of committee decision-making where private information is dispersed amongst members; the question how to model free-form communication remains open. Gerardi and Yariv (2007) make a counterpoint on this issue with a model that allows unrestricted deliberations before voting starts, with at least three voters and discrete types. They show that the set of outcomes that can be implemented with some communication procedure is the same for any voting rule except those requiring unanimity, so in a sense unrestricted pre-voting communication does not matter. Their construction relies on the fact that no single voter is pivotal, and hence no one has a unilateral incentive to deviate from any previously agreed upon outcome. However some of the equilibria supported by this argument may not survive free-form communications as they are not neologism-proof (Farrell, 1993). The authors also do not consider the possibility that free-form communications may allow all or the majority of voters to collectively deviate from equilibria they deem undesirable.

Our approach in Damiano, Li and Suen (2008b) is different. Instead of unrestricted costless communication, we model specific games that capture the phenomenon of strategic delay. Our model can be regarded as a bilateral signaling game: players who are informed about the correct decision signal their information by voting for that decision with a higher probability, even though such a vote may entail costly delay. In bargaining models with asymmetric information about the size of the pie to be divided (e.g., Cramton, 1992; Kennan and Wilson, 1993), signaling equilibrium exists in which the player who knows that the pie is relatively small is willing to endure a relatively long delay. The crucial difference between that class of models and ours is that delay is not a pure waste in our model. Costly delay helps improve the quality of the committee decision, and can produce a Pareto improvement to both committee members compared to an environment in which no signaling can take place.

In a follow-up study using a more general setup, Daminano, Li and Suen (2009) address the issue of what is the optimal deadline for reaching an agreement. It turns out answering this question is possible in a continuous time framework, where each uninformed committee member controls the flow probability of accepting the alternative. We show that the optimal deadline is deterministic; positive only if the ex ante probability of the disagreement state is neither too high nor too low; finite and increasing in the ex ante degree of conflict when it is positive. We plan to apply the same continuous time analysis to study other design issues such as temporary delay and super-majority rules in repeated committee voting.

7. Concluding Remarks

We have assumed throughout this article that there is common knowledge about committee members' preferences (their optimal rule of aggregating all the information) and competence (the quality of their private information). This assumption allows us to address the issues of providing incentives to gather information, dealing with strategic information aggregation and managing expertise within a unified framework. However, this may not be a good assumption for an ad-hoc committee formed to make a one-time decision. Further, for a committee that makes repeated decisions in similar environments, members may face constraints in the manipulation of their private information for the current decision due to concerns for their reputation,[20] and such concerns may affect their strategic interactions in making the current decision. Recently, a few authors have taken up the challenge of modeling reputational concerns in committee decision-making. Visser and Swank (2007) consider a committee of experts that have incomplete information about their own competence and care about how an outside evaluator updates about their competence levels from the decision that the committee makes. They assume that competent experts have identical information about the decision, and thus all experts have incentives to present a united front before the evaluator after the decision is made. In Levy (2007), committee members

[20] In Ottaviani and Sorensen (2006), agents are concerned with making recommendations that are validated ex post. Damiano, Li and Suen (2008a) study a model in which independent experts with reputational concerns for honesty and competence are constrained in their ability to distort their information, as people can make inferences by cross-checking those reports. There may be other constraints. For example, Milgrom and Roberts (1986) and Shin (1994) study the situation where agents' information can be concealed but cannot be distorted.

instead have private information about their own competence. This creates incentives for members to conform to each other in voting for the committee decision. Transparency about the votes helps by allowing the outside evaluator to update about the competence level of each member based on the votes as well as the decision.

In some collective decision-making environment, the committee that makes the decision for the group is formed endogenously by private choice, rather than exogenously fixed or chosen by an ultimate decision-maker. Direct democracy is a classical example, and so are elections in a representative democracy when participation is not mandatory. In Osborne, Rosenthal and Turner (2000), agents with different preferences and participation costs self-select themselves to join in committee decision-making. They show that the committee endogenously formed in an equilibrium is small and consists of members of extreme preferences. Their result suggests that there is too little participation. However, when agents have private information about the decision, participation of an agent can have positive as well as negative externalities on other agents: it brings to the committee information that benefits others, but reduces the probability that others are decisive. These two sides of costly participation are investigated in Ghosal and Lockwood (2004), and Borgers (2004).

Finally our discussion of committee decision-making is based on a model with fixed agenda. The committee is asked to make a binary decision regarding "the status quo" or the "alternative." Most actual decisions are richer than that, but there are also practical (as well as theoretical) problems with collective decision making when the menu of choices is unrestricted. In political science there is a rich literature on agenda setting and legislative bargaining, which deals with the formulation of policy proposals to ensure winning coalitions (Baron and Ferejohn, 1989). Much of this literature is developed in a framework with conflicting policy preferences but little or no role for information aggregation. Adding a policy proposal stage to our model of committee decision-making raises a number of interesting issues about the relationship between agenda-setting power and the strategic use and sharing of private information. It will be a fruitful area for research.

References

Aghion, Philippe and Jean Tirole (1997) 'Formal and Real Authority in Organizations,' *Journal of Political Economy* 105, pp. 1–29

Alonso, Ricardo and Niko Matouschek (2008) 'Optimal Delegation,' *Review of Economic Studies* 75, pp. 259-293

Aumanna, Robert and Sergio Hart (2003) 'Long Cheap Talk,' *Econometrica* 71, pp. 1619–1660

Austen-Smith, David and Jeffrey Banks (1996) 'Information Aggregation, Rationality and the Condorcet Jury Theorem,' *American Political Science Review* 90, pp. 34–45

Austen-Smith, David and Timothy Feddersen (2006) 'Deliberation, Preference Uncertainty and Voting Rules,' *American Political Science Review* 100, pp. 209–218

Baron, David and John Ferejohn (1989) 'Bargaining in Legislatures,' *American Political Science Review* 83, pp. 1181–1206

Bergemann, Dirk and Juuso Valimaki (1996) 'Learning and Strategic Pricing,' *Econometrica* 64, pp. 1125–1149

Borgers, Tilman (2004) 'Costly Voting,' *American Economic Review* 94, pp. 57–66

Blume, Andreas and Oliver Board (2007) 'Noisy Talk,' *Theoretical Economics* 2, pp. 395–440

Bolton, Patrick and Christopher Harris (1999) 'Strategic Experimentation,' *Econometrica* 67, pp. 349–374

Caillaud, Bernard and Jean Tirole (2007) 'Consensus Building How to Persuade a Group,' *American Economic Review* 97, pp. 1877–1900

Chan, Jimmy and Wing Suen (2007) 'Media as Watchdogs The Role of News Media in Electoral Competition,' University of Hong Kong working paper

Chwe, Michael (1999) 'Minority Voting Rights Can Maximize Majority Welfare,' *American Political Science Review* 93, pp. 85–97

Condorcet, M.J.A.N. de Caritat (1875) 'An Essay on the Application of Analysis to the Probability of Decisions Rendered by a Plurality of Voters,' abridged and translated in I. McLean and A.B. Urken, eds., *Classics of Social Choice*, 1995 (Ann Arbor: University of Michigan Press)

Cramton, Peter (1992) 'Strategic Delay in Bargaining with Two-Sided Uncertainty,' *Review of Economic Studies* 59, pp. 205–225

Crawford, Vincent and Joel Sobel (1982) 'Strategic Information Transmission,' *Econometrica* 50, pp. 1431–1451

Damiano, Ettore, Hao Li and Wing Suen (2008a) 'Credible Ratings,' *Theoretical Economics* 3, pp. 325–365

Damiano, Ettore, Hao Li and Wing Suen (2008b) 'Delay in Strategic Information Aggregation,' University of Toronto working paper

Damiano, Ettore, Hao Li and Wing Suen (2009) 'Optimal Deadlines for Agreements,' University of Toronto working paper

Davis, Michael (1994) ' The Value of Truth and the Optimal Standard of Proof in Legal Disputes,' *Journal of Law, Economics, and Organization* 10, pp. 343–359

DeGroot, Morris H. (1970) *Optimal Statistical Decisions* (New York: McGraw-Hill)

Dessein, Wouter (2002) 'Authority and Communication in Organizations,' *Review of Economic Studies* 69, pp. 811–838

Dessein, Wouter (2007) 'Why a Group Needs a Leader Decision-Making and Debate in Committees,' University of Chicago working paper

Dewatripont, Mathias and Jean Tirole (1999) 'Advocates,' *Journal of Political Economy* 107, pp. 1–39

Doraszelski, Ulrich, Dino Gerardi and Francesco Squintani (2003) 'Communication and Voting with Double-Sided Information,' *Contributions to Theoretical Economics* 3, Article 6

Downs, Anthony (1957) *An Economic Theory of Democracy* (New York: Harper Collins)

Duggan, John and Cesar Martinelli (2001) 'A Bayesian Model of Voting in Juries,' *Games and Economic Behavior* 37, pp. 259–294

Farmer, Amy and Dek Terrell (2001) 'Crime versus Justice Is there a Trade-off?' *Journal of Law and Economics* 44, pp. 345–366

Farrell, Joseph (1993) 'Meaning and Credibility in Cheap Talk Games,' *Games and Economic Behavior* 5, pp. 514–531

Feddersen, Timothy and Wolfgang Pesendorfer (1998) 'Convicting the Innocent The Inferiority of Unanimous Jury Verdicts under Strategic Voting," *American Political Science Review* 92, pp. 23–35

Feddersen, Timothy and Wolfgang Pesendorfer (1996) 'The Swing Voter's Curse,' *American Economic Review* 86, pp. 408–424

Feddersen, Timothy and Wolfgang Pesendorfer (1997) 'Voting Behavior and Information Aggregation in Elections with Private Information,' *Econometrica* 65, pp. 1029–1058

Gerardi, Dino and Leeat Yariv (2007) 'Deliberative Voting,' *Journal of Economic Theory* 134, pp. 317–338

Gerardi, Dino and Leeat Yariv (2008) 'Information Acquisition in Committees,' *Games and Economic Behavior* 62, pp. 436–459

Gersbach, Hans (1995) 'Information Efficiency and Majority Decisions,' *Social Choice and Welfare* 12, pp. 363–370

Gilligan, Thomas and Keith Krehbiel (1987) 'Collective Decision-Making and Standing Committees an Informational Rationale for Restrictive Amendment Procedures," *Journal of Law, Economics, and Organization* 3, pp. 287–335

Gilligan, Thomas and Keith Krehbiel (1989) 'Asymmetric Information and Legislative Rules with a Heterogeneous Committee,' *American Journal of Political Science* 33, pp. 459–490

Ghosal, Sayantan and Ben Lockwood (2004) 'Information Aggregation, Costly Voting and Common Values,' University of Warwick working paper

Goltsman, Maria, Johannes Horner, Gregory Pavlov and Franceso Squintani (forthcoming) 'Mediation, Arbitration and Negotiation,' *Journal of Economic Theory*

Holmstrom, Bengt (1983) 'On the Theory of Delegation,' in *Bayesian Models in Economic Theory*, 1983, ed. M. Boyer and R.E. Kilstrom (New York: North-Holland)

Keller, Godfrey, Sven Rady and Martin Cripps (2005) 'Strategic Experimentation with Exponential Bandits,' *Econometrica* 73, pp. 39–68

Kennan, John and Robert Wilson (1993) 'Bargaining with Private Information,' *Journal of Economic Literature* 31, pp. 45–104

Klevorick, Alvin K., Michael Rothschild and Christopher Winship (1984) 'Information Processing and Jury Decision-making,' *Journal of Public Economics* 23, pp. 245–278

Koriyama, Yukio and Balazs Szentes (2007) 'A Resurrection of the Condorcet Jury Theorem,' University of Chicago working paper

Krishna, Vijay and John Morgan (2001) 'A Model of Expertise," *Quarterly Journal of Economics* 116, pp. 747–775

Krishna, Vijay and John Morgan (2004) 'The Art of Conversation Eliciting Information from Experts through Multi-stage Communication,' *Journal of Economic Theory* 117, pp. 147–179

Kovac, Eugen and Tymofiy Mylovanov (2006) 'Stochastic Mechanisms in Settings without Monetary Transfers Regular Case,' University of Bonn working paper

Letterie, Wilko and Otto Swank (1997) 'Learning and Signaling by Advisor Selection,' *Public Choice* 92, pp. 353–367

Letterie, Wilko, Otto Swank, and Hendrik van Dalen (2000) 'When Policy Advisors Cannot Reach a Consensus,' *Social Choice and Welfare* 17, pp. 439-461

Li, Hao (2001) 'A Theory of Conservatism,' *Journal of Political Economy* 109, pp. 617–636

Li, Hao and Wei Li (2009) 'Optimal Delegation with Limited Commitment,' University of Toronto working paper

Li, Hao, Sherwin Rosen and Wing Suen (2001) 'Conflicts and Common Interests in Committees,' *American Economic Review* 91, pp. 1478–1497

Li, Hao and Wing Suen (2004) 'Delegating Decisions to Experts,' *Journal of Political Economy* 112, No. 1, Pt. 2, pp. S311–S335

Marino, Anthony (2007) 'Delegation versus Veto in Organizational Games of Strategic Communication,' *Journal of Public Economic Theory* 9, pp. 979–992

Martinelli, Cesar (2006) 'Would Rational Voters Acquire Costly Information?' *Journal of Economic Theory* 129, pp. 225–251

Martinelli, Cesar (2006) 'Rational Ignorance and Voting Behavior,' *International Journal of Game Theory* 35, pp. 315–335

Meirowitz, Adam (2007) 'In Defense of Exclusionary Deliberation Communication and Voting with Private Beliefs and Values,' *Journal of Theoretical Politics* 19, pp. 310–327

Melumad, Nahum and Toshiyuki Shibano (1991) 'Communication in Settings with No Transfers,' *Rand Journal of Economics* 22, pp. 173–198

Milgrom, Paul (1981) 'Good News and Bad News Representation Theorems and Applications,' *Bell Journal of Economics* 12, pp. 380–391

Milgrom, Paul and John Roberts (1986) 'Relying on Information of Interested Parties,' *Rand Journal of Economics* 17, pp. 350–391

Mukhopadhaya, Kaushik (2005) 'Jury Size and the Free Rider Problem,' *Journal of Law, Economics and Organization* 19, pp. 24–45

Myerson, Roger (1979) 'Incentive Compatibility and the Bargaining Problem,' *Econometrica* 47, pp. 61–73

Mylovanov, Timofiy (2006) 'Veto-based Delegation,' *Journal of Economic Theory* 138, pp. 297–307

Osborne, Martin, Jeffrey Rosenthal and Matthew Turner (2000) 'Meetings with Costly Participation,' *American Economic Review* 90, pp. 927–943

Ottaviani, Marco, and Peter Sorensen (2006) 'Professional Advice,' *Journal of Economic*

Theory 126, pp. 120–142

Persico, Nicola (2004) ‘Committee Design with Endogenous Information,” *Review of Economic Studies* 71, pp. 165–194

Sah, R. and J. Stiglitz (1986) ‘The Architecture of Economic Systems Hierarchies and Polyarchies,’ *American Economic Review* 76, pp. 716–727

Sah, Raaj and Joseph Stiglitz (1988) ‘Committees, Hierarchies and Polyarchies,’ *Economic Journal* 98, pp. 451–470

Shin, Hyun-song (1994) ‘The Burden of Proof in a Game of Persuasion,’ *Journal of Economic Theory* 64, pp. 253–264

Suen, Wing (2004) ‘The Self-Perpetuation of Biased Beliefs,’ *Economic Journal* 114, pp. 377–396

Szalay, Dezso (2005) ‘The Economics of Clear Advice and Extreme Options,’ *Review of Economic Studies* 72, pp. 1173–1198

Visser, Bauke and Otto Swank (2007) ‘On Committees of Experts,’ *Quarterly Journal of Economics* 122, pp. 337–372