Information Aggregation in Organizations

Inauguraldissertation
zur Erlangung des akademischen Grades
eines Doktors der Wirtschaftswissenschaften
an der Universität Mannheim

vorgelegt von

Elisabeth Schulte

Mannheim 2006

Referent: Prof. Dr. Hans Peter Grüner
Koreferent: Prof. Felix Kübler, Ph.D.
Dekan: Prof. Dr. Ulrich Falk

Tag der mündlichen Prüfung: 13.12.2006

Meinen Eltern

Acknowledgements

First of all, I would like to thank Hans Peter Grüner for his very valuable support. I am particularly grateful for his collaboration concerning Chapter 5 of this dissertation. I thank Felix Kübler for valuable comments. Many thanks go to my colleagues at the Lehrstuhl für Wirtschaftspolitik, especially to Thomas Gall, Alexandra Kiel and Kerstin Gerling. Chapter 2 benefitted a lot from the collaboration with Alexandra and Kerstin. I am particularly grateful for endless discussions with Thomas Gall and with Frank Rosar. Moreover, I thank Université Toulouse I for their hospitality and Emmanuelle Auriol, Michel LeBreton, Jaques Crémer and Jean Tirole for helpful discussions. Financial support from the Deutsche Forschungsgemeinschaft and through an EU Marie Curie fellowship is gratefully acknowledged.

Abstract

This dissertation contributes to the analysis of information aggregation procedures within organizations. Facing uncertainty about the consequences of a collective decision, information has to be aggregated before making a choice. Two main questions are addressed. Firstly, how well is an organization suited for the aggregation of decision-relevant information? Secondly, how should an organization be designed in order to aggregate information efficiently?

The main part (Chapters 2-4) deals with information aggregation in committees. A committee is a decision-making institution in which several individuals take part in the decision procedure and possibly hold private, decision-relevant information. In Chapter 2, a survey of the recent literature on committees is provided.

In Chapter 3, we study information aggregation in a committee whose members have heterogeneous preferences. Preference heterogeneity may interfere with information aggregation if the organization members disagree on how the information should be mapped into a decision. We study the performance of majority voting as a mechanism to aggregate information, when agents have the possibility to publicly announce their information before voting takes place. We identify conditions under which full information aggregation is possible.

In Chapter 4, we compare the performance of two alternative decision procedures facing partially conflicting interests among decision makers. The two decision procedures differ with respect to the extent to which they allow communication among decision makers. We find that limiting the individuals' access to communication may enhance decision quality.

In Chapter 5, we depart from the committee framework. Decision-relevant information is no longer private, but centrally available. The question is how to efficiently organize the information aggregation procedure. The organization is evaluated in terms of two dimensions, speed and quality of decision making. We assume that it takes time to read information, and that agents make a mistake with a certain probability when carrying out a processing task. The extent of parallel information processing affects the time it takes to reach a decision. The quality of the decision is affected by processing imperfections.

Contents

Chapter 1

Introduction

In many organizations, collective decision procedures involve more than one individual, but less than all the individuals who are affected by the decision. In most developed countries, decisions which affect the whole population are made by an elected government. In firms, important decisions are made by the board of directors. Hiring decisions are made by hiring committees, monetary policy is conducted by a monetary policy committee and criminal trial juries decide upon the fate of a defendant.

Why are collective decisions delegated to small groups instead of involving all individuals or delegating to a single decision maker? We explore an informational explanation for this question. If the consequences of a collective decision are uncertain, uncertainty can be reduced by acquiring and processing information. Consider for instance a decision whether to implement a reform or not, and suppose that the reform will be welfare-enhancing in some states of the economy, but detrimental for welfare in others. The better the decision-making institution is endowed with information which allows to predict the state of the economy more accurately, the better it will be able to identify the welfare-maximizing decision.

By involving more than a single decision maker, the collective decision can potentially be based on more information if the decision makers have private information which is otherwise not accessible. On the other hand, even if all the decision-relevant information is centrally available, there may be a role for the participation of several agents in the decision making procedure because this allows for the delegation of information processing tasks.

This dissertation elaborates on both perspectives towards information aggregation, the aggregation of private information in collective decision making procedures, and the aggregation of centrally available information in a decentralized manner through the delegation of processing tasks. Two major questions are addressed: Is it possible to use the collective resources, i.e. individuals' private information and information processing capability, efficiently? Are some forms of organization better suited to use these resources than others?

In Chapters 2-4, we consider the case in which decision-relevant information is spread among organization members, and study the performance of voting mechanisms as a means to aggregate this information. In Chapter 5, we study the case in which information is centrally available, but it needs to be aggregated in order to understand it. There, the focus is on the optimal aggregation procedure in the presence of information processing imperfections.

Chapter 2 originates from joint work with Kerstin Gerling, Hans Peter Grüner and Alexandra Kiel which has been published in the European Journal of Political Economy (21) in 2005. We provide an overview over the recent game-theoretical literature on decision making in committees. It is assumed that committee members possess private, decision-relevant information. The questions which are addressed by this literature are twofold: (i) How well is the political process (e.g. majority voting) suited to aggregate private information, and (ii) How should the political process be optimally designed in the presence of private information? These questions are studied in different settings, with exogenous or endogenous information, with or without conflicting interests among committee members, and with or without the possibility to communicate. Among other insights, the literature suggests that information aggregation is severely limited if committee members have conflicting interests, even if they are allowed to exchange views prior to casting their votes.

However, this conclusion might be too pessimistic, as it relies crucially on the specific informational environment which is assumed, namely private information being soft. That is, agents can always claim to possess information which favors their preferred decision, which causes a limit to information aggregation. In Chapter 3, we suggest an alternative approach, assuming that information is verifiable. Committee members are allowed to make public speeches, but they cannot pretend to have information which they in fact do no have. After the communication stage, a decision is made by majority voting. We analyze different settings, and derive conditions under which full information aggregation is possible. In particular, if preferences are common knowledge and each committee member is endowed with information full information aggregation is possible despite preference heterogeneity.

In Chapter 4, we compare two decision procedures, which differ with respect to the extent to which decision makers can exchange information. In a setting in which the full information aggregation result of Chapter 3 does not apply, we allow agents with partially conflicting interests to reveal private, decision-relevant information either to the entire group of decision makers (open debate), or only to decision makers with the same interests (group debate). If information disclosure is observed by all decision makers, agents may strategically withhold information. As a consequence, individual votes will be less sensitive to observable information. We show that limiting the individuals' access to communication may yield a higher expected social surplus.

Chapter 5 is joint work with Hans Peter Grüner, which has been accepted for publication in the Journal of Economic Theory. We study a decision problem in which all the decision-relevant information is centrally available, but it needs to be aggregated in order to identify the optimal decision. The information aggregation technology captures two important aspects of human processing ability. First of all, we assume that it takes time to read an information item. Secondly, we introduce information processing imperfections. Thus, the evaluation of a decision making organization is carried out in terms of two dimensions, speed and quality of decision making. We consider the task to choose an alternative out of a given set of alternatives. An agent, when choosing between two options, makes a mistake with a certain (exogenous) probability, and chooses the worse option. We derive the hierarchical organization in which the best option is chosen with highest probability.

The thesis is organized in such a way that the chapters can be read independently of each other. As the chapters are closely connected, this involves a certain degree of repetition for readers who read through the entire manuscript at the benefit of readers who read the chapters selectively. All references are collected in the bibliography.

Chapter 2

Information Acquisition and Decision Making in Committees: A Survey

2.1 Introduction

How do committees work? And how should they be designed? A recent game-theoretic literature has added useful insights to the theory of committee decision making. The role of this chapter is to provide an overview over the recent developments in this field and to relate it to some current debates on the design of committees for international decision making.

It is not the aim to provide a comprehensive study of all contributions in this field, but rather to deliver an introduction to this branch of literature, to relate different results and to unsheathe important assumptions concerning the study of committees. The focus is on the strategic behavior of committee members facing decisions that are made only once and that are irreversible. For a survey of the literature which deals with repeated strategic interaction and political institutions, see Piketty (1999).

The formal study of committees is old. In his classical contribution Condorcet (1785) described a committee as a mechanism that efficiently aggregates information. In his famous jury theorem he argues that (i) increasing the number of informed committee members raises the probability that an appropriate decision is made, and (ii) the probability of making the appropriate decision will converge to one as the number of committee members goes to infinity.

It is useful to relate the modern literature on committees to Condorcet's early insight. Condorcet's analysis was based on a simple setup where (i) individuals always base their choice on their information, (ii) individuals obtain their information at zero cost, (iii) all individuals have the same objective: to make a correct decision, and (iv) individuals do not exchange views before voting. In many cases of interest some (or even all) of these assumptions do not hold. Some voting rules may induce individuals not to vote in accordance with their own information. When information acquisition is costly, individuals provide less effort in large committees. Conflicting interests may lead to the misrepresentation of information. And communication may affect individual voting behavior when information is distributed asymmetrically. Recent papers have therefore addressed the issue of committee decision making when one or more of Condorcet's assumptions do not hold.

We first give an introduction to the model setup applied by most of the papers we survey. In Section 2.3, we briefly turn to the Condorcet Jury Theorem in this setup and unsheathe the (explicit and implicit) assumptions. We then discuss contributions which study the role of strategic voting in committees (Section 2.4). Next we look at papers that analyze incentives for information acquisition (Section 2.5). We then turn to the role of differences in preferences (Section 2.6) and after that to pre-vote communication (Section 2.7). Section 2.8 summarizes some experimental results on committee decisions and relates them to the theoretical literature.

Key questions related to committee design are "How large should a committee be?", "Who should be in a committee?", "What is the optimal decision rule?", and "What should the delegating body do with the committee's decision?". We will summarize and compare the answers to these questions in Section 2.9.

2.2 The model setup

In this section, we introduce the model setup applied by most of the papers we survey. Throughout the chapter, we will refer to this section.

2.2.1 The basic model

A policy decision $x \in X$ has to be made by a group of n agents, called the committee. This decision is determined by voting. We restrict attention to the case where the policy decision is binary, i.e. $X = \{a, b\}$. Examples are decisions concerning the implementation of a reform or to maintain the status quo, or the acquittal or conviction of a defendant.

Each agent i casts a vote $v_i \in \{a, b\}$ for one of the alternatives. The voting rule is characterized by a number of votes $k \in \{1, ...n\}$ needed to implement alternative a:

$$x(k, v) = \begin{cases} a, & \text{if } \sum_{i=1}^{n} 1_{v_i = a} \geq k \\ b, & \text{otherwise.} \end{cases}$$

where $1_{v_i=a} = 1$ if i votes for a and 0 otherwise. That is, for $k = \frac{n+1}{2}$, alternative a is implemented if at least a simple majority of the committee members votes in favor of it, otherwise b is implemented. For $k = n$, the committee needs a unanimous agreement in order to implement alternative a.

There exist two possible states of nature $s \in \{A, B\}$. The utility agents derive from the implementation of a policy is state-dependent. For example, the utility from convicting a defendant depends on whether the defendant is guilty or innocent. We assume that all agents prefer policy $x = a$ in state of the world $s = A$ and $x = b$ in state of the world $s = B$. Unfortunately, the state of the world is not known. Agents share the common belief $\pi = prob\{s = A\}$, $0 < \pi < 1$. In the basic version of the model we assume that agents' preferences can be represented by the following von Neumann-Morgenstern utility function:

$$u(x, s) = \begin{cases} 1, & \text{if } x = s \\ 0, & \text{otherwise.} \end{cases}$$

This function captures the following properties: agents suffer symmetrically from possible wrong decisions and also value correct decisions symmetrically. Thus, the agents' common objective is to implement the appropriate policy,[1] which amounts to guessing the state of the world. Prior to the voting decisions,

[1] When agents have homogeneous preferences, the terms "correct" or "appropriate" policy refer to the one that matches the state of the world. When agents' preferences differ, it is not straightforward to define what the appropriate policy is. In that case, we need further assumptions regarding the agents' preferences to justify the desirability of a certain outcome.

each agent i receives an imperfect signal $\sigma_i \in \{\alpha, \beta\}$ with the following stochastic properties:

Assumption 2.1

$$\begin{aligned} prob\{\sigma_i = \alpha \,|s = A\} &= prob\{\sigma_j = \alpha \,|s = A\} := q_0 \quad \forall i, j, \\ prob\{\sigma_i = \beta \,|s = B\} &= prob\{\sigma_j = \beta \,|s = B\} := q_1 \quad \forall i, j, \end{aligned}$$

and $\frac{1}{2} < q_0, q_1 < 1$.[2]

Let $prob\{s = A\,|\sigma_i = \alpha\} := q_\alpha \ \forall i,\ prob\{s = 1\,|\sigma_i = 1\} := q_\beta \ \forall i$, and assume:

Assumption 2.2

$$\tfrac{1}{2} < q_\alpha, q_\beta < 1.$$

The parameters π, q_0 and q_1 describe the informational environment. We will refer to q_α and q_β as the quality or the accuracy of the signal.

The timing is as follows:

1. Nature chooses the state of the world according to the prior belief π.

2. Nature sends a signal to each agent according to the rules specified above.

3. Agents simultaneously vote for one of the alternatives.

4. The policy is implemented according to the voting rule k and payoffs realize.

In order to capture other important aspects of committee decision making, such as conflicting interests, costly information acquisition or the possibility to exchange information before voting, this framework has to be enriched, as specified below.

[2] We restrict attention to the case of equally skilled decision makers. However, the distribution and quality of decision skills matter for committee voting. Two papers that have dealt with this issue are Nitzan and Paroush (1982) and Karotkin and Nitzan (1995). Nitzan and Paroush (1982) show that with different decision skills the optimal decision rule is a voting rule that grants more weight to better qualified decision makers.

2.2.2 Conflicting interests

Agents may have different preferences over the possible outcomes of the game. Let $u_i(x,s) \geq 0$ denote the utility agent i derives from the implementation of policy x in state of the world s. We assume that $u_i(a,A) \geq u_i(a,B)$, and $u_i(b,B) \geq u_i(b,A)$, which in a jury setup means that agent i does not derive higher utility from the acquittal of a defendant if the defendant is guilty rather than innocent nor from the conviction of an innocent instead of a guilty.[3] More generally speaking: An agent should derive at least as much utility from the implementation of an "appropriate" policy as from the same policy, implemented in a state of the world in which it is not appropriate.

In the case of heterogeneous preferences, it is not straightforward what the terms "appropriate" or "correct" policy mean, nor which outcome of the game can be valued as desirable. We use these terms in the following sense: For each state of the world, there exists a policy that would be agreed upon by the (majority of) agents if the state of the world was known for sure. It is desirable to choose this "correct" policy with high probability. If we impose certain symmetry assumptions on agents' preferences, this desirability criterion is equivalent to maximizing Bentham's welfare function.

Maximizing expected utility, agent i prefers policy a over policy b if

$$p(\cdot)\, u_i(a,A) + (1-p(\cdot))\, u_i(a,B) \geq p(\cdot)\, u_i(b,A) + (1-p(\cdot))\, u_i(b,B),$$

where $p(\cdot)$ denotes i's assessment of the probability that the state of the world is A. We have

$$\begin{aligned} p(\cdot)\, u_i(a,A) + (1-p(\cdot))\, u_i(a,B) &\geq p(\cdot)\, u_i(b,A) + (1-p(\cdot))\, u_i(b,B) \Leftrightarrow \\ p(\cdot) &\geq \tfrac{u_i(b,B)-u_i(a,B)}{u_i(a,A)+u_i(b,B)-u_i(a,B)-u_i(b,A)}. \end{aligned}$$

Agent i's preferences can be represented by $\theta_i = \frac{u_i(b,B)-u_i(a,B)}{u_i(a,A)+u_i(1,1)-u_i(a,B)-u_i(b,A)}$. Following (among others) Feddersen and Pesendorfer (1998), we will refer to θ_i as i's "threshold of reasonable doubt". Agent i prefers alternative a to be implemented as long as he assesses the probability that state A realized to be higher than his threshold of reasonable doubt.

[3] In the course of the analysis, we will often refer to the jury interpretation of the information aggregation game in order to simplify the exposition. This is not to say that we believe conflicting interests are especially present in jury trials.

The assumptions stated above ensure that the denominator is positive. The nominator is negative, if $u_i(a, B) > u_i(b, B)$, which means that i prefers to acquit the defendant even if he was sure that he is guilty. θ_i is greater than 1, if $u_i(b, A) > u_i(a, A)$, which implies that i prefers to convict the defendant even if he was sure that he is innocent.

If an agent with preferences satisfying either $u_i(a, B) \geq u_i(b, B)$ or $u_i(b, A) \geq u_i(a, A)$ exists, we call him partisan. All non-partisan voters' preferences are characterized by $\theta_i \in (0, 1)$.

2.2.3 Information acquisition

If the agents do not receive their signals for free, but each of them has to invest effort in order to learn something about the state of the world, this will be modeled by adding another stage to the game in which prior to receiving signals, agents simultaneously decide about their investment in effort. We will be more specific about the investment technology in the course of the discussion of papers which deal with this aspect.

The timing in a model with an information acquisition stage is then as follows:

1. Nature chooses the state of the world according to the prior belief π.
2. Agents simultaneously make their investment decisions.
3. Dependent on the agents' investment decisions, nature sends signals according to the information technology.
4. Agents simultaneously vote for one of the alternatives.
5. The policy is implemented according to the voting rule k and payoffs realize.

2.2.4 Communication

In order to capture the possibility of communication within our model framework, a cheap talk stage prior to the voting stage is introduced to the game. An agent's strategy in this setting is a plan what to report depending on the signal he receives and for which alternative to vote, depending on the signal he got, the report he made and the messages he received. As we shall see, the introduction of a cheap talk stage may considerably complicate the analysis.

The timing in a model with a communication stage is as follows:

1. Nature chooses the state of the world according to the prior belief π.
2. Nature sends signals to the agents.
3. Agents exchange messages.
4. Agents simultaneously vote for one of the alternatives.
5. The policy is implemented according to the voting rule k and payoffs realize.

2.3 The Condorcet Jury Theorem in a nutshell

In this section, we briefly outline the Condorcet Jury Theorem (hereafter CJT). The model framework is as described in Section 2.2.1. The first part of the CJT (CJT I) states that the probability of making an appropriate decision in a committee which decides via simple majority voting is increasing in the number of informed committee members. The second part (CJT II) states that this probability converges to one as the number of informed committee members goes to infinity. Committee members are informed in the sense that every agent is endowed with a signal which is correlated with the true state of the world as specified in Section 2.2.1.

We give an example to illustrate the mechanics of the model. Let $\pi = \frac{1}{2}$, and let $q_0 = q_1 = q$. So we have $q_\alpha = q_\beta = q$.

Condorcet assumed that each committee member votes according to the signal he received. Such a voting behavior is called *informative voting*.[4] The probability that a single committee member makes a correct decision is thus q.

Now consider a committee with three members. The probability that the majority receives a correct signal (and thus votes in favor of the correct policy) is $\binom{3}{2} q^2 (1-q) + q^3 = 3q^2 - 2q^3$.

The probability to make a correct decision is higher in the committee consist-

[4]See e.g. Austen-Smith and Banks (1996).

ing of three persons than if one agent determines the policy alone, iff

$$\begin{aligned} 3q^2 - 2q^3 &> q \Leftrightarrow \\ \left(3q - 2q^2 - 1\right) q &> 0 \Leftrightarrow \\ (2q-1)(1-q) q &> 0 \Leftrightarrow \\ \tfrac{1}{2} &< q < 1. \end{aligned}$$

CJT II rests on the law of large numbers: The probability that a fraction q of committee members receives correct signals converges to one. Since $q > \frac{1}{2}$, the probability that at least a fraction $\frac{1}{2}$ receives correct signals is converging to one, too.

The explicit and implicit assumptions on which Condorcet's argumentation rests are the following:

1. Agents vote according to their information (informative voting).
2. Agents are endowed with information at no cost.
3. Agents' interests are perfectly aligned.
4. Agents do not exchange their information before voting.

Actually, the third and the fourth assumption do not play a role if the first one is satisfied. Note that the first assumption implies that the mapping of information into the decision – and therewith the decision quality – is exogenously given. As we shall see, conflicting interests and the possibility to communicate may play a major role for the strategic interaction if this mapping is endogenous. In the following, the assumptions will be relaxed. We start with dropping the first one and allow agents to determine their voting choices by themselves, i.e. we allow for strategic voting.

2.4 Strategic voting

Unsatisfied with the strict behavioral assumptions underlying the CJT, a branch of the voting literature investigates the features of strategic voting from a game-theoretic point of view. In fact, restricting the strategy space to informative voting is not innocuous, since rational voters will generally not vote informatively in equilibrium. The reason is that a committee member tends to neglect his own information, while he tries to deduce other committee members' private information from their voting behavior. Then, he might either not vote any longer according to his own private information or even abstain, if he feels less informed and shares common values with the electorate.

2.4.1 Rational voting

Austen-Smith and Banks (1996) were among the first to unsheathe the implicit behavioral assumptions that individuals vote *sincerely*, i.e. as a member of a collective, each individual selects the alternative he would have selected when voting alone, and *informatively*, i.e. each committee member's decision reflects the signal he received before.

The authors start from a simple Bayesian game as described in Section 2.2.1. A strategy for a player in this game is a plan which vote to cast for each signal he may get. Informative voting requires voting for a after having observed α and voting b after having observed β. Committee members share homogeneous preferences for selecting the better of the two alternatives. The decision is taken by majority vote without abstentions. When making their decision, agents take into account the common prior probability π in favor of state A and a private, imperfect signal which they received about the true state of the world. Austen-Smith and Banks show that it is the structure of individuals' information that endogenously generates heterogeneous policy preferences. Based on being pivotal, a rational voter deduces other individuals' private signals and by incorporating this additional equilibrium information into his beliefs, he tends to neglect his own private information. It follows that informative voting by all individuals is generally not rational.

We provide an example to illustrate this result. Consider a three person committee in an informational environment $\{\pi, q_0, q_1\}$ as defined in Section 2.2.1. To check whether informative voting by all agents constitutes an equilibrium strategy profile, we need to determine the best response of a player i to the other players voting informatively. This player needs to worry only about the case in which he is pivotal, i.e. in case one player votes in favor of alternative a and the other one in favor of alternative b, otherwise his voting choice will have no effect on the final outcome and on his utility. Given that the other players vote

informatively i can infer their signals. He then updates his beliefs about which state of the world realized according to all the information available to him.
If $\sigma_i = \alpha$, i assesses the probability that the state of the world A realized to be

$$prob\left\{s = A\left|\sigma = (\alpha, \alpha, \beta)\right.\right\} = \frac{\pi q_0^2\left(1-q_0\right)}{\pi q_0^2\left(1-q_0\right) + \left(1-\pi\right)\left(1-q_1\right)^2 q_1}.$$

If $\sigma_i = \beta$, i assesses the probability that the state of the world 0 realized to be

$$prob\left\{s = A\left|\sigma = (\alpha, \beta, \beta)\right.\right\} = \frac{\pi q_0\left(1-q_0\right)^2}{\pi q_0\left(1-q_0\right)^2 + \left(1-\pi\right)\left(1-q_1\right) q_1^2}.$$

His best response is to vote for alternative a iff $prob\left\{s = A\left|\cdot\right.\right\} > \frac{1}{2}$ and for alternative b otherwise. Thus, informative voting is rational iff the following two conditions hold:

$$\begin{aligned}\frac{\pi q_0^2\left(1-q_0\right)}{\pi q_0^2\left(1-q_0\right) + \left(1-\pi\right)\left(1-q_1\right)^2 q_1} &> \frac{1}{2} \\ \frac{\pi q_0\left(1-q_0\right)^2}{\pi q_0\left(1-q_0\right)^2 + \left(1-\pi\right)\left(1-q_1\right) q_1^2} &\leq \frac{1}{2}.\end{aligned}$$

It should be obvious that these condition hold for some, but not for all parameter constellations. For example, for $\pi = .5, q_0 = .6$ and $q_1 = .9$, we have $\frac{\pi q_0^2(1-q_0)}{\pi q_0^2(1-q_0)+(1-\pi)(1-q_1)^2 q_1} > \frac{1}{2}$ and $\frac{\pi q_0(1-q_0)^2}{\pi q_0(1-q_0)^2+(1-\pi)(1-q_1)q_1^2} > \frac{1}{2}$. That is, in this informational environment, i's best response to the other players' informative voting strategies is to vote for alternative a independently of his signal. Thus, informative voting by all players fails to be an equilibrium strategy profile. In fact, there is just one exceptional case for which Condorcet's behavioral assumption is met, namely if the majority rule k is the optimal method of aggregating individuals' private information. An aggregation rule is optimal, if and only if, the rule being used, it is rational for each member to vote informatively if all others do so. In a second step, the authors allow for variations in the structure of individuals' information. Specifically, if the decision makers receive two independent private signals or additionally a public signal, the paper's main result applies again.

But then, it is no longer assured that majorities invariably do better than individuals in selecting the better of two alternatives. It is easy to verify that in the example above the game has an equilibrium in which all players always vote for alternative a. If we expected this equilibrium strategy profile to be played, it would be better to delegate the decision to only one agent for whom (this is also easy to verify) informative voting would be rational and the probability of making a correct decision would be higher.

It has to be mentioned that voting games generally have multiple equilibria. It was not the aim of Austen-Smith and Banks to identify all of them but to show that informative voting by all agents may fail to be one of them.

As McLennan (1998) argues, this does not imply that the CJT fails to hold when allowing for strategic voting. Instead, whenever there is a profile of votes for which, given the voting rule k, the Jury Theorem holds there must also be a profile of votes which achieves the same and which is a Nash equilibrium of the voting game. This is due to the fact that the voting game played by agents with homogeneous preferences constitutes a common interest game in which the maximizer of the common utility function is also a best response for an agent with the same objective function.

Given the voting rule, an equilibrium in which all players vote informatively may not exist, and thus the efficient use of the available information may be precluded. The information environment crucially affects the outcome of the game induced by the voting rule. Therefore the appropriateness of the use of a majority rule hinges upon the characteristics of the encountered situation. In order to make use of all the information available to the electorate, i.e. to induce all voters to vote informatively in equilibrium, the voting rule has to be adjusted to the information environment.

2.4.2 Abstention

In contrast to previous models of voter turnout, which traditionally focus on the costs and benefits of voting, Feddersen and Pesendorfer (1996, 1999b) present an informational explanation for the existence of abstention and roll-off.

Their model considers the behavior of voters with heterogeneous preferences in a voting game as described in Section 2.2.1. There are three types of voters: two types of partisans, who, regardless of the state of the world, either prefer alternative a or b, and independents, who prefer to select the option that matches the true state of the world. The quality of signals differs across agents: While some agents get a useless signal, others receive a perfect signal. These informed agents are certain about the realization of the state variable, which affects the utility of all voters.

Applying insights from the theory of auctions, the authors show that with private information and common values less informed voters have an incentive to abstain rather than to vote for either alternative even though voting is costless and though all abstainers strictly prefer one alternative over the other. In fact, the uninformed independent voters' reason to cast a vote is to compensate for the partisans. That is how they maximize the probability that the informed voters decide the election. Having achieved this compensation, it is optimal to delegate the decision via abstention to more informed voters.[5] An implication of Feddersen and Pesendorfer's findings is that differences in information about the different items on the ballot will make voters abstain on some issues and vote on others. Hence, the authors also provide an explanation for the existence of roll-off. Moreover, they go on to show that even though significant abstention occurs in large elections, the outcome of the election is almost always the same as with perfect information.[6]

2.4.3 Unanimity

With the minimization of criminal trials' expected wrongful verdict costs being a common social aim, unanimous jury verdicts were usually seen as a means to reduce the probability of convicting an innocent while increasing the probability of acquitting a guilty defendant (see e.g. Klaven and Zeisel (1966) or Adler (1994)). Feddersen and Pesendorfer (1998, 1999a) were the first to challenge this basic intuition by taking into account strategic voting by jurors.

They consider a simple voting game as described in Section 2.2.1, in which the alternatives a, b are to be interpreted as the acquittal or conviction of a defendant and the states of the world A, B denote innocence and guilt of the defendant respectively. Jurors are homogeneous with respect to their preferences, characterized by a threshold of reasonable doubt $\theta \in (0, 1)$ as introduced in Section 2.2.2. A juror prefers the defendant to be sentenced if he believes in the defendant's guilt with a probability higher than the threshold of reasonable doubt.

Each juror behaves as if his vote was pivotal. Under the unanimity rule, this is the case if all other jurors agree, which reveals additional information about the

[5] Coupé and Noury (2002) provide some empirical support for this 'swing voter's curse'.

[6] Fey and Kim (2002) elaborate a correct proof of the first proposition, which does not require any alteration of the paper's results.

true state. Such information may overwhelm the juror's private assessment of the case and cause him to vote with the others. As a result and in opposition to the outcome under naive voting, even in a large jury, the probability of convicting an innocent defendant must stay bounded away from zero. The information aggregation potential of a unanimous voting mechanism is bounded.

The rationale causing this result is as follows: In any equilibrium of the voting game in which the defendant is sometimes convicted and sometimes acquitted, there must be a jury member who votes to acquit with positive probability, but not for sure. For this to be equilibrium behavior it must be that the event of being pivotal does not contain overwhelming evidence of guilt. Hence, the other agents must vote to convict with high enough probability even if they receive a signal indicating innocence. Next, note that in the case of a unanimity requirement for conviction, the difference between the event that any voter i is pivotal and the event of a conviction is just i's single vote. Thus, the probability that a convicted defendant is innocent as well as the probability of convicting an innocent defendant are bounded away from zero independent of jury size.

The authors also draw comparisons between the unanimity rule and a wide variety of special majority verdicts of a size less than unanimity, including simple majority rule. Given any voting rule requiring a fixed fraction of votes to condemn, Feddersen and Pesendorfer are able to explicitly solve for the corresponding unique symmetric, responsive Bayesian equilibrium. It turns out that among those voting rules, unanimous jury verdicts may be least appropriate to track the truth and result in higher probabilities of both kinds of error, i.e. convicting the innocent and acquitting the guilty. More precisely, a jury theorem holds for all voting rules other than unanimity. While with an increasing size of the jury, the probability of making a mistaken judgement goes to zero for all voting rules, except for unanimity, even the opposite may be the case for unanimity: the probability of convicting an innocent defendant may even increase with the size of the jury. Based on an example confronting different voting rules for a fixed jury size, the authors finally conclude that in order to reduce the probability of convicting an innocent defendant, any other super-majority rule with a large jury is more appropriate than unanimity.

However, the authors admit that the degree to which strategic voting and private information matter in actual juries is crucial for the ultimate outcome and thus at last emerges to be an empirical question. Martinelli (2002b) explains the inferiority of unanimity by exploiting the analogy between unanimity rule and a single prize auction. Specifically, in a very general setting, which include continuous or discrete signals as special cases, a classical paper by Milgrom (1979) obtains a necessary and sufficient condition for the winning bid to converge to the true price of an object in a large auction. Martinelli shows that a similar condition is necessary for full information aggregation in the limit, and that under this condition the probability of convicting an innocent converges to zero. The condition is that for every state such that some voter would like to acquit, there must be signals that make the voter arbitrarily sure of being in that state rather than in a state in which every voter would like to convict. With a discrete number of signals and two states, the proposed condition cannot be satisfied if there are no perfectly informative signals.

Coughlan (2000), on the other hand, argues in favor of unanimity voting by extending the basic setting to include more realistic features of actual jury trials. In particular, if a jury faces the risk of a mistrial, i.e. a certain number k of votes $\left(> \frac{n}{2}\right)$ is required for both, conviction as well as acquittal, unanimity voting outperforms majority voting. The reason is that voters have a greater incentive to vote informatively, because (i) the event of being pivotal does not shape beliefs systematically into one direction and (ii) with an innocent (guilty) signal a voter assesses the event to be pivotal for acquittal more likely than for conviction (vice versa). Coughlan shows that given informative voting, the probability to acquit the guilty as well as the probability to convict the innocent decrease in k under mistrial voting.

2.5 Incentives for information acquisition

The analysis of strategic voting points out that committees need not yield better results than individuals. However, this result is based upon fixed decision rules which may no longer be appropriate in larger committees. If decision rules are adjusted properly, then the Condorcet Jury Theorem still holds. Reasons for limits on committee size can be found if one drops the rather harsh assumption that information comes for free. Obviously, costly information acquisition in a committee constitutes a public good – it tends to be under-provided. Increasing the number of committee members reduces the incentives for information acquisition, because the probability that an individual member is pivotal declines.

The formal analysis of these incentives yields different results. Papers in which information acquisition is a discrete choice come to the conclusion that larger committees may yield poorer decisions and a lower social surplus. On the other hand, when information acquisition is continuous, larger committees may still yield better informed decision making in the aggregate despite lower individual incentives for information acquisition. After considering these two versions concerning the information investment technology and the conditions under which the optimal outcome can be reinstalled, we turn to information acquisition incentives that arise from the relation between the committee and a parent body with limited authorization to revise the committee's decision. Finally, we briefly discuss the effect committee composition may have on the incentives to acquire information.

2.5.1 Information acquisition: continuous choice

Martinelli (2002a) shows that a large committee may anticipate the true state of the world with probability close to one, although the committee members do not know anything about the state of the world ex ante and information acquisition is costly. This is true if the information cost function satisfies certain conditions. The paper's main point is that "rational ignorance" on the part of committee members is consistent with a well-informed committee in the sense of forecasting the correct state of the world with a high probability.

The model framework Martinelli applies is as described in Section 2.2.3. When an agent invests in information, he receives a signal whose accuracy depends on the information investment. More precisely, if a voter buys x units of information, he receives a signal which is correct with probability $1/2 + x$. The information investment costs follow a strictly convex, twice differentiable function.

Both states of the world are equally likely ex ante and payoffs for the implementation of the correct policy are symmetric for the independents. Moreover, there are partisans among the voters, who, regardless of the state, either prefer a or b. The ex ante probability of being a type a partisan is equal to being one of type b.

The timing is as follows: First, nature selects the voters' types which are private information. Second, voters decide simultaneously about the quality of their information. Information investments are not observable. Then, voters can either vote for a or b. Majority wins.

Partisans never acquire information and always vote for their preferred alternative. Because of the symmetric structure, there is no problem with insincere voting among the moderate voters under majority rule (see Austen-Smith and Banks, 1996).

If marginal information cost at zero information is positive and the number of voters is large enough, there exists no equilibrium with information acquisition. If marginal cost is zero at the point of zero information, there exists a unique equilibrium in which all independent voters acquire the same amount of information, in turn depending on its marginal benefit, and in which they vote sincerely.

If the second derivative of the information cost function is also zero at the point of zero information, the probability of choosing the right alternative (in the sense that alternative X is chosen if the state of the world is X) converges to one as the number of voters goes to infinity. If it is positive, but bounded, the probability of choosing the right alternative converges to some value between 1/2 and 1, depending on the parameters. If the value of the second derivative converges to infinity as information converges to zero, success probability converges to 1/2 as the number of voters goes to infinity. Moreover, elections become very close as the electorate grows.

The intuition for these results is the following: If information is cheap enough (first and second derivative being zero at zero information), independent voters will acquire some, because they are pivotal with positive probability. The existence of partisans and the imposed symmetry ensure that the probability of being pivotal does not fall too fast as the number of voters grows. Although information acquisition of the individual moderate voter goes to zero as the number of voters goes to infinity, it does so slowly enough to allow for the effect of large numbers to kick in. The poor information of the individual voter does also explain why elections are close in this setup. Moreover, this feeds back on information acquisition incentives, since in close elections the probability of being pivotal is high.

2.5.2 Information acquisition: binary choice

For the case of binary information investment decisions, Mukhopadhaya (2003), Ben-Yashar and Nitzan (2001) and Persico (2004) show that – due to a free rider problem in information acquisition – a larger jury may make worse decisions. This is well in line with Martinelli's (2002) result since marginal information investment costs at the point of zero information are positive in the case of a binary choice variable.

A nice and intuitive example is Mukhopadhaya's (2003) two player game with a perfect signal. Each player may purchase a perfect signal about the state of the world. Both players share information. The game has two asymmetric pure strategies Nash equilibria where one of the two jury members acquires the perfect signal. The other one is a mixed strategies equilibrium in which both players buy the signal with a positive probability. In the case with one single decision maker, the decision maker always decides to buy the signal. Hence, in the mixed strategies equilibrium the probability of making a correct decision is lower than in the one decision maker case.

Mukhopadhaya's game with an imperfect signal is as follows: First nature chooses the true state of the world. Then each agent may decide to invest in the costly signal. Next, all the jurors pool their information which is possible because they have a common objective. In the vote they all agree on the decision that has to be taken.

The game has a symmetric mixed strategies equilibrium. The author shows that for extreme (high) values of the signal's precision, one juror is more likely to reach a correct decision than three jurors. The author provides an example where the probability of making a correct decision is first increasing and then monotonously decreasing in the number of jurors.

In a similar setting, Persico (2004) determines the optimal voting mechanism consisting of the voting rule and the committee size. A voting mechanism has to aggregate information efficiently as well as to provide proper incentives to acquire information. The underlying questions are: Under which circumstances should majority determine collective decisions, when is it better to rely on more stringent measures of consensus? And how large should a group of decision makers be?

Persico designs the optimal voting mechanism by choosing the number of committee members n and the plurality rule k needed to implement policy b. The optimal mechanism maximizes agents' expected utility from the collective decision. Costs enter only insofar as a smaller committee is chosen only if this does not decrease expected utility. To solve the induced game, Persico restricts attention to pure and monotone strategies equilibria – more precisely, he is only interested in the most efficient equilibrium in pure and monotone strategies. Voters are homogeneous and can aggregate information only through their votes, communication is considered to be impossible.

Plurality is determined à la Austen-Smith and Banks (1996), i.e. k is chosen such that the maximum number of agents vote informatively in equilibrium. The basic trade-off is that by enlarging the committee (combined with an adjustment of the voting rule k), the decision becomes more accurate, but voters become less pivotal such that their information acquisition incentive shrinks. If the committee is too large, no one will acquire information since the probability of deciding the final outcome is too low. Thus, there exists a bound on n.

There exists a bound on k, too: The optimal fraction of votes k/n needed to implement policy b can never be greater than approximately the accuracy of the signal. This is true irrespective of agents' preferences, i.e. independent of the agents' threshold of reasonable doubt θ (see Section 2.2.2). This implies that large pluralities (in the extreme unanimity) are optimal only if the information available to committee members is sufficiently accurate.

As n grows, the optimal decision rule k converges to simple majority. The optimal number of committee members is non-increasing in information acquisition cost and ceteris paribus larger if (i) the preference structure is more symmetric, i.e. the threshold of reasonable doubt is closer to $1/2$, (ii) the prior is more diffuse and (iii) the signal is less accurate.

The above statement, that an increase in committee size decreases the incentives to acquire information, is only one part of the story. The opposite is true as long as the optimal plurality rule k/n converges to signal accuracy: Because the optimal decision rule k itself moves with n (and at half speed of the growth of n, as Persico shows), it creates an opposite-directed effect. An increase in n associated with a minor increase in k makes the individual voter in fact more pivotal than in the smaller committee, as long as k/n has not yet reached approximately signal quality. As k/n converges beyond signal quality to simple majority, both effects operate in the same direction and information acquisition incentives indeed decrease with further increases in committee size.

Persico shows that his results also hold for heterogeneous agents by considering two possible types. If agents' types are observable and preferences sufficiently diverging, it is optimal to leave the decision to a group of only one type, using the optimal rule that would be used if these agents were the only ones the mechanism designer is interested in. There exists no voting rule which incorporates votes of both types such that all types vote sincerely. If agents' types are unobservable but the number of agents of each type is common knowledge, again the decision rule of only one group is used, modified in a way such that the votes of the other type are "sterilized", i.e. in equilibrium the other type votes in order to correct the voting rule regardless of their signal and the type whose decision rule is used votes sincerely.

The author admits that the restriction to pure strategies may be critical, since allowing some agents playing mixed strategies might indeed lead to superior outcomes, because pure strategy players have stronger incentives to acquire information. Moreover, the role of communication is entirely neglected.

2.5.3 Relation to the delegating authority

If there is a delegating authority, incentives to acquire information depend upon the relation between the committee and the delegating authority. Gilligan and Krehbiel (1987) argue that restricting the ability of the parent body to amend committee proposals may enhance the informational role of committees. The model of Gilligan and Krehbiel (1987) works as follows. There are two players, the floor and the committee. A policy x has to be chosen, x is a real number. The desired policy of the committee is larger than the floor's desired policy, however the desired policy of both actors is affected in the same way by a shock. The committee members may acquire costly information about this shock. The committee reports a bill, then the floor makes a decision. Under the unrestricted procedure the floor may pick any policy after obtaining the report, under the restricted policy it may either accept the bill or stick to the status quo. Gilligan and Krehbiel show that restrictions on the ability of the parent body to amend committee proposals may provide better incentives to acquire information and may lead to an outcome which is better both for the parent body and the committee.

2.5.4 Committee composition and incentives

A recent literature (Dewatripont and Tirole, 1999, and Beniers and Swank, 2004) studies the impact of the composition of committees on incentives for information acquisition. Dewatripont and Tirole study a setup where a principal uses monetary incentives to induce two agents to search for information of opposing content. Beniers and Swank study a similar setup without monetary transfers. In this setup, it may be useful to delegate information acquisition to two individuals with biased - and opposing - interests. This provides them with incentives to acquire information. Delegation to individuals with unbiased preferences is optimal when costs for information acquisition are low.

2.6 Conflicting interests

Different approaches have recently been developed to study the role of conflicting interests in committees. In some jury models jurors care differently about wrongful acceptance and wrongful rejection of a hypothesis, i.e. conflicting interests are modeled in terms of differing thresholds of reasonable doubt as described in Section 2.2.2. It may even be the case that committee members are not completely aware of their own preferences. Cai (2001) proposes a model in which committee members learn their preferences by investing costly effort and in Grüner and Kiel (2004), there is a model in which committee members' preferences are interdependent but agents are only aware of their own private value component. Finally, Li et al. (2001) have a model that allows for communication and shows how the degree of conflicts affects the incentives to exaggerate reports about the own private information and hence the efficiency of the committee's decision.

2.6.1 Partially conflicting interests

Feddersen and Pesendorfer (1997) analyze the performance of elections with heterogeneous voters when there is uncertainty about a one-dimensional state variable. Despite heterogeneity and a vanishing fraction of informatively voting agents, large elections perform well. The authors show that the information environment is crucial in determining the effectiveness of elections as information aggregation mechanisms.

The model is similar to the one in Feddersen and Pesendorfer (1998), as introduced in Section 2.4.3, except that the threshold of reasonable doubt is heterogeneous across agents and constitutes an agent's private information. We briefly resume the setup:

A population of n voters has, by a given majority rule k/n, to decide between two alternatives, a or b, e.g. in the jury interpretation acquittal or conviction of the defendant. Alternative b is elected if it receives at least a fraction k/n of the votes. Each voter's payoff depends on his specific preference type (the threshold of reasonable doubt), the true state of nature (innocence or guilt) and the chosen alternative. Preference types are drawn independently from a commonly known distribution. Each voter knows his own preference type but does not know the other voters' types. Every voter receives a private signal which is correlated with the true state of nature. Jurors vote strategically, i.e. they condition their vote on the event of being pivotal, taking into account the other agents' strategies. They are not allowed to communicate before voting. Timing is as follows: Nature draws the agents' types according to a commonly known distribution and the state of the world according to the commonly known prior. Agents learn their type and their private signal, and vote. In order to solve the game, the authors restrict attention to symmetric Bayesian Nash equilibria in which players do not use weakly dominated strategies.

In equilibrium, preference types can be divided into three groups: those who always vote for either alternative and those who vote informatively, i.e. condition their vote on their private signal.

As the size of the electorate goes to infinity, the fraction of players who condition their votes on their private information goes to zero. Nevertheless, voting fully aggregates information in the sense that with probability close to one the alternative is elected that would have been chosen if all private information was common knowledge (the k/n-median's preferred outcome). Moreover, with probability close to one, in equilibrium the winning alternative receives a fraction of votes close to the fraction required to win.

The intuition for the results is the following: Voters condition their voting strategy on the event of being pivotal. This implies that beliefs about the state of nature concentrate on the state in which, given the equilibrium strategy profile, it is most likely that alternative b receives a fraction k/n of votes. As the voting population grows to infinity, this evidence becomes very strong and the fraction of voters for whom the own signal is decisive goes to zero. Voters behave as if they for sure were in the predicted state. Voters with low thresholds of doubt are very concerned about acquitting the guilty. The information they can infer out of their signal and the event of being pivotal does not convince them of the defendant's innocence beyond their threshold of doubt, so they vote to convict him. A similar argument holds for voters with high thresholds. They are so concerned about convicting the innocent, that equilibrium information does not convince them of the defendant's guilt. Medium types are convinced by their signal and use it for their decision. Although there is only a vanishing fraction of voters who condition their votes on information, the number of them goes to infinity. Since these are the voters who determine the outcome of the election, the election performs well in the limit.

Gerardi (2000) complements Feddersen and Pesendorfer's (1997) analysis by comparing unanimous and non-unanimous voting rules in the same model setup. Moreover, his analysis complements the results in Feddersen and Pesendorfer (1998) by taking preference heterogeneity into account. He shows that any non-unanimous decision rule is asymptotically efficient. In large committees, the unanimous rule almost never leads to the decision for which unanimity is required.

A symmetric Bayesian Nash equilibrium exists for any decision rule and any jury size. Under unanimity, the probability that an innocent is convicted converges to zero as the jury size grows to infinity, but the probability to acquit the guilty converges to one. Thus, protecting the innocent comes at the prize of acquitting the guilty. Moreover, the probability that a convicted defendant is innocent converges to zero. Under any non-unanimous rule, which is defined as a fraction k/n of voters required to convict the defendant, the probability to convict the innocent as well as the probability to acquit the guilty converge to zero.

Gerardi does not propose an optimal voting mechanism, his main points are that unanimity is not optimal in large juries, whereas non-unanimous rules are asymptotically efficient. The first result is not surprising, since the existence of a single voter with sufficiently extreme preferences suffices to free the defendant. As the jury size converges to infinity, the existence of such a voter will be very likely, even though evidence of guilt, on its part inferred out of being pivotal, becomes stronger.

On the other hand, non-unanimous rules, characterized by a fraction of votes required to convict the defendant, have different asymptotic properties. Take any super-majority rule. As the jury size becomes larger, the number of acquit-votes needed to acquit the defendant grows and, at the same time, the interval of types always voting to acquit shrinks, since the evidence of the defendant's guilt becomes stronger. So, the asymptotic efficiency of the non-unanimous rule is not surprising either.

To sum up, Gerardi's results strengthen the findings of Feddersen and Pesendorfer (1998, 1999a) about the inefficiency of unanimity.

2.6.2 Preference uncertainty and committee size

Taking costly information acquisition into account, Cai (2001) develops a model of committee size when agents are uncertain about the state of the world and in addition about their own policy preferences. The state of the world is a point on the real line on which the payoff from the implemented policy (again a point on the real line) depends. When exerting nonverifiable effort, an agent learns his policy preferences and receives a noisy signal about the state of the world. Thus, conflicting interests among committee members arise from information acquisition in this model.

N agents are selected into a committee by a principal who represents society's preferences in the sense that his policy preferences coincide with an uninformed agent's expected preferences. The committee members' task is to acquire costly information and to report it to the principal who uses it to update his beliefs about the state of the world and then decides upon the policy variable. Neither information acquisition decisions nor signals are observable. Updated beliefs and incentive compatibility conditions for committee members constitute the elements for a Perfect Bayesian Nash Equilibrium of this multiple stage incomplete information game for a given committee size N.

Since information is soft and informed committee members know their policy preference which may differ from the principal's, there exist incentives for strategic information manipulation.

The game is solved by backward induction. First, the author characterizes a reporting equilibrium of the information aggregation stage. Attention is restricted to strictly monotonic (reversible) reporting strategies. In this equilibrium, committee members convey all their information (except their policy preferences) to the principal. Committee members with no policy biases report truthfully, and those with policy biases exaggerate by a multiple of their policy preference. Uninformed committee members prefer not to submit any information at all, because their expected policy preferences coincide with the principal's and a signal announcement would only create additional noise.

The principal makes his decision as the (to exaggeration adjusted) mean of all reports. The author shows that this equilibrium is essentially the unique equilibrium consisting of reversible reporting strategies, essentially unique in the sense that all equilibria in fully reversible strategies have identical outcomes. Moreover, he proves that it is the most efficient among all equilibria. Given this reporting equilibrium, the committee members' incentives to gather information and the optimal committee size are studied.

Information acquisition incentives limit the size of the committee. The optimal number of committee members is lower than the first best, i.e. if no incentive problems would exist and any expected gain from additional information is traded off against participation costs. Heterogeneity of preferences plays the crucial role: If preferences were identical among all agents and information acquisition costs did not exceed participation costs, the first best would be attainable since interests would be completely aligned.

Interestingly, information acquisition incentives and therewith committee size may increase in preference heterogeneity. Benefits from information acquisition contain two elements: firstly, the policy becomes more informative and secondly, the agent learns his policy preferences and gets the chance to manipulate the policy in his own favor. Information acquisition serves as an insurance facing preference uncertainty. When information acquisition costs are high, such that the committee size is limited by members' shirking tendency, an increase in preference heterogeneity raises the value of this insurance and mitigates the shirking problem.

2.6.3 Decision rules with interdependent preferences

Grüner and Kiel (2004) analyze collective decision problems in which individual bliss points are correlated but not identical. The authors compare the performance of two specific decision mechanisms as regards different degrees of correlation.

All agents obtain private information about their most desired policy. However, the individually preferred decision of a group member does not only depend on his own private information but also on the other group members' private information. The decision problem is characterized by a parameter which measures the extent to which private information affects all individuals. This specification includes the private values case for the lower bound of the interdependency parameter and the common values case for its upper bound.

Participation in the decision is not voluntary and monetary transfers are excluded a priori. Instead, the mechanisms maps the profile of individual announcements into the collective decision. Attention is restricted to two specific decision mechanisms, the median and the mean mechanism. The main difference between these two mechanisms is how they weight individual announcements of private information. Under the median mechanism changes in extreme positions are disregarded, since the median alone determines the final decision. On the contrary, the nature of the mean mechanism is to take all available information into account. Therefore, under the mean mechanism extreme positions influence the decision.

The main result of this paper is the identification of two symmetric Bayesian Nash equilibria of the respective games. The performance of the mechanisms depends upon the extent to which spill-over effects affect the economy. With weak interdependencies, the median mechanism dominates the mean mechanism, whereas with strong interdependencies it is better to use the average as decision mechanism.

The reason is that for weak interdependencies the equilibrium strategy under the median mechanism implies announcement behavior close to truth-telling whereas the equilibrium strategy under the mean mechanism leads to strong exaggeration of private information. Therefore, average taking is outperformed by ignoring some of the information. Since the degree to which interdependencies influence untruthful announcement behavior is stronger under the mean mechanism, this intuition holds for a wide range of interdependencies, only for very high degrees it is reversed.

There are two main points that remain unconsidered in this paper. First, the authors abstract from individual rationality considerations. It should be noted that even individuals who do not participate in the mechanism would be affected by the collective decision. Allowing for voluntary participation would require to take into account endogenous participation constraints. Another question is the design of an optimal mechanism for the class of collective decision problems studied.

2.6.4 Partially conflicting interests and information sharing

Li, Rosen, and Suen (2001) analyze committee decisions when members have partially conflicting interests and possess private information. The decision variable is binary and private information is a continuous variable. Interests are partially conflicting in the sense that there is an area of disagreement for some values of information, i.e. an area for which one voter prefers alternative a but the other one prefers b. Preferences are common knowledge. The paper's main result is that information cannot be fully shared and voting procedures arise as the equilibrium method of information aggregation.

Since information is private, committee decisions are made on the basis of the members' reports of their private data. The authors show that under these circumstances, information cannot be fully shared among committee members in the sense that it is not possible to exactly deduce private signals from reports. Efficient (or full) sharing requires that the committee decision responds to small changes in any members' data. This property fails in any Bayesian equilibrium of any decision-making procedure. Incentive compatibility implies that continuous data observed by each member are partitioned and transformed into rank order categories. In equilibrium, personal thresholds are chosen to undo the presumed biases of other committee members, but not by enough to completely nullify the information of the others. The coarsening of information balances incentives to exaggerate information and incentives to share information. Nevertheless, incentives for manipulation and counter-manipulation generate a larger area of disagreement among committee members than is implied by their inherent conflicts in preferences.

It is shown that the greater the latent consensus among members, the greater are the opportunities for presenting private data in finer categories. On the other hand, conflicting interests among committee members impose an upper bound on how fine information partitions can be. Indeed, the quality of committee decisions improves with the degree of consensus.

The authors demonstrate that delegation to one single member is Pareto dominated by committee decision-making. When the committee rule is chosen appropriately, gains from sharing information outweigh distortions from information manipulation regardless of the extent of conflicts in the committee. The reason is that committee members are more cautious when casting the decisive vote than if they were to make the decision alone in order to take advantage of the other members' data. Moreover, if one member is known to have data of higher quality, the others cast their decisive votes less frequently. The coarsening of information implies that the committee decision rule is ex post inefficient.[7]

[7]Another model in which voters receive signals from a continuous distribution, but do not

2.7 Communication

Most of the models we have discussed so far rely upon Condorcet's original assumption that individuals in a committee do not communicate before they cast their votes. One might argue that this whole branch of literature bears useless results when extensive discussion and exchange of information precedes the votes, hence reducing incentives for strategic voting à la Austen-Smith and Banks (1996). Indeed, in models with homogeneous preferences it is the restriction of the information transmission via votes which causes the inefficient usage of the information. However, as we have seen in the model of Li, Rosen and Suen (2001), jury members might only make very limited use of the possibility of information transmission when conflicts of interest are present. Hence, studying pre-vote communication seems a worthwhile task.

Recent theoretical work indeed proves that pre-vote communication gives rise to a new kind of equilibrium, since first-stage mutual pre-vote revelation of impressions about the true state of the world replaces the need for committee members to augment private information by deducing other members' information from their voting behavior. Then, in the voting stage, all voters - provided the absence of preference heterogeneity - agree on the preferred alternative: the one that matches the state of the world with highest probability given the signal vector. This voting behavior in the second stage is obviously an equilibrium. Moreover, given the second stage voting behavior and given truthful revelation, it is optimal to reveal the own information truthfully in the first stage.

In this section, we review the (small) literature on pre-vote communication in collective decision making, resulting in significant alterations of the results of previously discussed committee voting models. After having presented how Gerardi and Yariv (2003a) derive the irrelevance of voting rules for the equilibrium outcome ensued from pre-vote communication and costless information, we further follow them to demonstrate the need for a limited optimal committee size when information becomes costly (Gerardi and Yariv, 2003b). Moreover, Doraszelski et al. (2003) show that pre-vote exchange of views improves the decision in the presence of conflicting interests.

---

have the option to communicate before voting, was developed by Duggan and Martinelli (2001). They show that in this setting the probability of making a wrong decision converges to zero as the jury size goes to infinity for every voting rule other than unanimity rule.

2.7.1 Irrelevance of voting rules

Gerardi and Yariv (2003a) consider a voting game where committee members may communicate before they cast their votes. The role of communication is studied in a setup where individuals obtain their information without investing effort, as described in Section 2.2.4. In the first stage, each individual obtains a signal which is his private information. His utility is a function of the collective action – which may again take two values – and of the vector of signals. Preferences are homogeneous. In the second stage, the cheap talk stage, individuals may communicate. In the third stage, individuals cast their votes. The voting rule maps the vector of votes into the set of outcomes. The voting rule is characterized by the number of votes k required to implement a certain decision. The main result is that with cheap talk, the set of equilibrium outcomes is independent of the voting rule. As we have seen above, this is not the case in models without cheap talk.

2.7.2 Imperfect aggregation and incentives

Gerardi and Yariv (2003b) introduce a stage of costly information acquisition. Each individual may purchase a signal which is correlated with the true state of the world. Moreover, in this setup they are interested in the optimal information aggregation mechanism. Thus, the game contains three stages. In the first stage, the designer chooses an extended mechanism which consists of the size of the committee, the voting rule and a rule which specifies how the players can exchange messages before voting. In the second stage, the agents observe the mechanism and decide whether to purchase the signal. These choices are made simultaneously. In the third stage, members of the committee communicate according to the pre-specified rules and vote.

The main results are the following. First, the authors show that the optimal committee is of bounded size – as before. Second, they show that imperfect aggregation of the available information may yield a higher overall expected utility level than perfect aggregation. The reason is that imperfect aggregation induces more players to acquire information. The positive incentive effect may dominate the negative effect of wasting some information.

2.7.3 Communication and conflicting interests

Doraszelski et al. (2003) also analyze a voting game which includes a communication stage prior to voting. They consider a two person committee which has to decide whether to change the status quo or not. The status quo is optimal in one state of the world, changing it in the other one. Agents differ with respect to their threshold of reasonable doubt. Hence, there are conflicting interests in the committee. Preferences are the agents' private information.

The game proceeds as follows. Nature draws the agents' types according to a commonly known distribution. Agents receive a private signal which is correlated with the true state. In the one sender version of the game, one of the agents sends a message to the other one, in the two sender game both agents simultaneously send messages to each other. Both agents vote simultaneously and the collective decision is implemented. The status quo is maintained unless both vote in favor of changing it.

The authors look for Perfect Bayesian Equilibria with the additional requirement that agents do not use weakly dominated strategies. In order to explore the interaction of communication and voting, they restrict attention to responsive robust cutoff equilibria, i.e. equilibrium strategies which imply that the receiver and the sender condition their votes on their types, their signals and the sender's message. The authors provide a complete characterization of these equilibria and show that communication is beneficial. The main purpose of communication is to serve as a "double-check": If a player's information conflicts with his preferences, he uses the submitted information of the other player to confirm his own.

In a one sender version of the game in which agents possess information of different quality, the identity of the sender is irrelevant. The game in which the better-informed agent is the sender is outcome-equivalent to the game in which he is the receiver. After observing the signal the better-informed sender conditions the message decision on an event that has the same probability as the event on which the better-informed receiver conditions the voting decision after observing the message and the signal. The authors conclude that communication and voting are perfect substitutes in the sense that information not transmitted in the communication phase will be aggregated in the voting stage.

The authors compute ex ante expected utilities for the pure voting game, the one sender and the two sender game for a uniform type distribution. Agents' ex ante utility is larger with two senders than with one sender. And it is larger with one sender than with no sender. However, this is due to decreasing returns to scale.

Since adding communication complicates the analysis considerably, the authors concentrate on a framework with two players, two states and two signals. So, there is scope for future research to generalize this setup. For example, adding players would allow to compare different voting rules.

Ottaviani and Sorensen (2001) analyze the effects of the order of statements on the effectiveness of information transmission when speakers care about their reputation of expertise, but not about the decision. They show that sequential statements may lead to herding behavior and that expected decision quality may decrease when information quality of a committee member increases. However, the sequential statement mechanism can outperform simultaneous announcing if the prior is sufficiently unbalanced.

2.8 Experimental results

Guarnaschelli, McKelvey and Palfrey (2000) provide a first experimental study of jury decision making. The experimental setup is identical to the one described in models of voting under imperfect information such as the ones by Austen-Smith and Banks (1996) or Feddersen and Pesendorfer (1998, 1999a). Individuals were informed about the existence of a true state of nature and they were given a signal with known stochastic properties. This setup is used to test for the presence of strategic voting. It turns out that individuals indeed vote strategically.

In addition, the authors conducted a so-called straw poll experiment. In this experiment, individuals were asked for the simultaneous announcement of a signal after which the vote takes place. As argued in Section 2.7, there is an equilibrium of this game where all individuals announce their signal correctly in the first round and then cast identical votes in the second round. It turned out that over 90 % of the individuals revealed their signal in the first round. In the second round, about 84 % or more voters voted in a way consistent with the outcome of the first round. Voting with a straw poll improves results in one respect: In the experiment, the probability of one type of error was reduced while the other was not affected.

There are other results in the paper by Guarnaschelli et al. which are at odds with theoretical predictions. For example, inconsistent with the theoretical findings of Feddersen and Pesendorfer (1998), in an experiment with unanimity rule, the probability of convicting an innocent defendant declines with jury size.

The authors also find that voting behavior might be consistent with jury members playing a mixed strategies equilibrium. Actually, the Austen-Smith and Banks (1996) game has mixed strategies equilibria. Whereas some individuals vote sincerely, others always vote to convict and others mix. It should be noted that in this experiment payoffs from correct predictions and from incorrect ones are the same for all individuals. Individuals received 50 US cents if the group decision was correct and five cents if it was incorrect.

Blinder and Morgan (2000) present experimental research on the quality of group decision making. The setup of their basic experiment is the following: Individuals (or groups of individuals) are confronted with the result of a binary lottery. The lottery has an initial probability of 50 % for both possible states. The individuals were told that the probabilities would change at some randomly selected point of time during the experiment. The change would either go up to 70 % or down to 30 % for each of the events. The individuals were asked for a guess of the point of time when the random process was changed. They were punished for a decision that has been made too late and they received a reward for a decision that was correct.

Group members were presented exactly the same information as an individual decision maker: the outcome of the random draw. The first hypothesis that has been tested was that groups make decisions more slowly than single individuals. However, the hypothesis that the decision takes as much time in a group as individual decision making could not be rejected. The hypothesis that groups outperform individuals was instead strongly supported by the experimental data. The third hypothesis concerned the speed of group decisions with different decision rules. The authors compared majority and unanimity rule. It turned out that unanimity rule worked faster than majority rule. Moreover, under majority rule individuals tended to reach unanimous agreements.

The experimental results by Blinder and Morgan are very useful because they show that a group may indeed perform better than individuals in analyzing the same information set (the series of random draws). The fact that decisions do not seem to take more time when a group decides under unanimity rule is somewhat surprising. Further theoretical and experimental research on the issue of timing might help to clarify this issue.

2.9 Summary of the theoretical results

Actually, it is needless to say that the applicability of the theoretical results to the real world depends on the specific features of the encountered situation. As one will realistically expect to face e.g. a combination of the Condorcet Jury Theorem's underlying assumptions to be violated to various degrees or further complications, the results derived in this literature can only serve as a rough, but sound guideline.

Nevertheless, the review of this literature allows drawing some general conclusions on the optimal organization of committees. These insights will guide our discussion of actual committee decision making concerning the optimal size, composition and decision rule.

2.9.1 How large should a committee be?

The optimal size of a committee depends upon several issues. First, it plays an important role whether individual information is exogenous or endogenous. In Condorcet's ideal world information is exogenous, i.e. individuals do not need to invest costly effort in order to obtain information. Moreover, all agents share a common objective. In such a setup, larger committees always lead to better results.

Austen-Smith and Banks (1996) confirm this result when the majority rule is adjusted properly to the information environment. Insincere voting is not a problem in such situations and information is aggregated efficiently. As McLennan (1998) shows, if a "naive" version of the CJT in this setup holds, there always exists an equilibrium of the voting game played by strategically acting agents with the same properties. However, efficient information aggregation requires to adopt the decision rule à la Austen-Smith and Banks (1996).

When information acquisition is costly, the optimal committee size is finite. This result is derived in Mukhopadhaya (2003), Ben-Yashar and Nitzan (2001), and Persico (2004). According to Persico (2004), the optimal number of committee members is non-increasing in information acquisition costs and ceteris paribus larger if (i) the threshold of reasonable doubt is closer to $1/2$, (ii) the prior is more diffuse, and (iii) the signals are less accurate.

In a setup with conflicting interests, Li, Rosen and Suen (2001) demonstrate that delegation to one single member is Pareto dominated by committee decision-making. When the committee rule is chosen appropriately, gains from sharing information outweigh distortions from information manipulation regardless of the extent of conflicts in the committee.

According to Cai (2001), when information acquisition is costly and unobservable and agents may become aware of conflicting preferences, the optimal committee size is always smaller than in the absence of incentive problems. The reason is that noisy reports by committee members with policy biases are not as informative as if there were no incentives to distort information.

2.9.2 Who should be in a committee?

According to Li, Rosen and Suen (2001), committee decisions are better the more similar the committee members' preferences. With increasing conflicts of interest, the area of disagreement becomes larger. This, in turn, leads to stronger incentives for strategic manipulation of private information and thus to a decrease in expected decision quality.

Contrary to this result, Cai (2001) shows that – once there are conflicting interests among committee members – total social surplus may increase in preference heterogeneity. This can happen when moral hazard problems in information gathering severely limit the feasible committee size. Recall that information acquisition has two components in Cai's model: firstly, agents receive a signal about the state of the world and secondly, they become aware of their own preferences. In this model, two opposing effects are present: Heterogeneity in preferences has a direct negative effect on total social surplus, because it increases the noisiness of information on which the collective decision is based. This effect is comparable to the one in Li, Rosen and Suen (2001). On the other hand, heterogeneity here provides additional incentives to gather information, since agents can manipulate policy in their ideal way. Thus, when the positive participation effect dominates the negative noisiness effect, increasing heterogeneity can increase total social surplus.

With regard to the quality of signals, committee decisions improve with the quality of individual information. If one committee member has access to data of higher quality the other in the setting used by Li, Rosen and Suen (2001), then the other member takes advantage of the improved information by changing his voting threshold in order to defer the decision to the informed member. Thus, better-informed members are decisive more often. When communication takes place, it is irrelevant whether it is the sender or the receiver of information who has data of higher quality (Doraszelski et al. (2003)).

2.9.3 What is the optimal decision rule?

This survey of the game-theoretic literature on committee decision making has made clear that the decision rule has to be adapted to the specific problem at hand. Austen-Smith and Banks (1996) have shown that in order to make efficient use of the available information, the majority rule has to be adjusted to the distribution of signals and the initial prior distribution of states of the world. However, as the number of jury members grows, the optimal decision rule converges to simple majority (see Persico, 2004).

According to Feddersen and Pesendorfer (1998) and Gerardi (2000), unanimity and the absence of communication lead to biased and socially undesired decisions in large committees. If voters' preferences are homogeneous, the probability to convict an innocent stays bounded away from zero independently of jury size. When jurors differ with respect to their thresholds of reasonable doubt, the probability that an innocent is convicted with a unanimous verdict converges to zero as jury size grows to infinity, but the probability to acquit the guilty converges to one. So, protecting the innocent comes at the prize of acquitting the guilty. Under any non-unanimous rule, the probability to convict the innocent as well as the probability to acquit the guilty converge to zero.

The introduction of pre-vote communication among committee members alters some of these results. Gerardi and Yariv (2003a) show that if pre-vote communication is introduced into the basic model (see Sections 2.2.1 and 2.2.4), the voting rule becomes unimportant. Doraszelski, Gerardi and Squintani (2003) show that even if agents have partially conflicting interests, voting outcomes improve when pre-vote communication is allowed. Moreover, it is irrelevant if it is the better or the worse informed agent who sends a message prior to voting. The authors conclude that pre-vote communication and voting are perfect substitutes in the sense that the information not transmitted in the communication phase will be aggregated in the voting stage.

Gerardi and Yariv (2003b) show that imperfect aggregation of the available information may yield a higher overall expected utility level than perfect aggregation. The reason is that imperfect aggregation induces more players to acquire information. The positive effect of this may dominate the negative effect of wasting some information. However, imperfect aggregation mechanisms face a time-consistency problem. At the stage when the decision is to be made, committee members may agree not to stick to the decision procedure and to aggregate information efficiently. The use of imperfect mechanisms necessitates the ability to stick to the procedure.

Committee design has also been studied without paying attention the role of strategic voting, i.e. under the assumption that agents to vote sincerely. In some situations the optimal decision rule induces a behavior where everybody indeed votes sincerely. In those cases results from the literature on optimal decision rules under sincere voting also become relevant in settings where individuals have the option to vote insincerely. Major contributions on the optimal institutional design under sincere voting are Sah and Stiglitz (1988) Ben-Yashar and Nitzan (1997), Ben-Yashar and Paroush (2001), and Koh (1994, 2004a,b).

Sah and Stiglitz (1988) analyze different majority rules in a setting with symmetric abilities of committee members to identify the true state of the world. Moreover, they discuss the role of sequential voting in what they call a hierarchy. In such a hierarchy a project is accepted only if it is accepted on all hierarchy layers. Ben-Yashar and Nitzan (1997) have generalized their analysis by admitting different abilities, different payoffs from type 1, and type 2 errors, different priors, and state-dependent decision skills. They derive a general voting rule which weights votes according to these parameters. Ben-Yashar and Paroush (2001) provide a further generalization with respect to the choice among more than two alternatives. In a similar setting Koh (2004a) studies the optimal amount of information acquired when information acquisition is costly. Koh (1994) analyzes optimal decision rules with more than two possible realizations of individual signals. These rules may be interpreted as evaluation standards.

Chapter 3

Information Aggregation and Preference Heterogeneity in Committees

3.1 Introduction

Decision-making entities are often comprised of agents who represent different interests. The most obvious example of such a decision-making institution is the government in a representative democracy. If the consequences of the decision are uncertain, the quality of the decision benefits from exchanging information prior to making a choice. However, in case committee members' interests are not completely aligned, we cannot take information exchange for granted. Should we worry about preference heterogeneity interfering with information aggregation? This is the question the chapter is concerned with.

The idea that committee members pool private information relevant for the decision, and therewith make use of a broader information base than a single decision maker could access, dates back at least to Condorcet (1785). The advantage of involving a higher number of informed agents in the decision process is intuitive if the committee indeed makes use of the individual members' private information. However, if committee members' interests are not completely aligned, this cannot be taken for granted.

A recent game-theoretic literature has shown that we cannot take efficient information aggregation for granted even if preferences are perfectly aligned (e.g. Austen-Smith and Banks, 1996; for a survey see Chapter 2). The reason is that the individual voter cares only about his vote when it is pivotal. Obviously, there exist equilibria in which no single vote ever affects the outcome and voters do not use their information. Austen-Smith and Banks (1996) show (in a voting setting without communication) that the exploitation of all available information is generally not possible in equilibrium. But Condorcet's jury theorem may still apply to strategically acting agents with completely aligned preferences (McLennan, 1998), since they play a common interest game.

Conflicting interests among committee members may limit their ability to pool their information efficiently. When preferences are heterogeneous, it is not straightforward to decide how decision quality should be measured. One such measure used in the literature (e.g. Feddersen and Pesendorfer 1997, 1998; Gerardi 2000) is the extent of information aggregation, i.e. the probability with which the collective decision would be the same if all the information was common knowledge. In this paper, we will also use this benchmark. Full information aggregation is desirable if committee members agree on which decision to make if the state of the world is known. Then, preferences are heterogeneous in the sense that voters differ with respect to their 'thresholds of doubt', i.e. with respect to how convinced a voter must be that a certain alternative is the correct choice in order to support that alternative.

In voting games without communication, full information aggregation requires that private information is transmitted via individual votes. If preferences are too heterogeneous, then full information aggregation via majority voting is impossible because beliefs concentrate around the threshold of doubt of the politically decisive voter, the median preference type. Voters whose thresholds are too far away from the median do not condition their votes on information (see Feddersen and Pesendorfer 1997).

On the other hand, committee members – at least in small committees – generally have access to another instrument to pool their information other than individual votes: they may exchange views prior to making a decision. In this regard, the information aggregation potential of committees is still not well understood. Most papers restrict attention to voting and neglect the role of communication within the committee. Exceptions are Coughlan (2000), Doraszelski et al. (2003), Gerardi and Yariv (2003a, b), and Austen-Smith and Feddersen (2002). Coughlan (2000) and Gerardi and Yariv (2003b) deal with committees composed of agents with homogeneous preferences for which complete information aggregation is possible because agents share a common goal.

The papers by Doraszelski et al. (2003) and Austen-Smith and Feddersen (2002) indicate that in committees with heterogeneous preferences, information aggregation is severely limited even if pre-vote communication is allowed. In these papers, information is soft and preferences are assumed to be private information. In such a setting, agents with extreme preferences always make statements which favor their preferred decision. Hence, in equilibrium, the information content of a statement is limited. Austen-Smith (1990a,b) studies information transmission in an agenda-setting game where preferences are common knowledge and information is soft. Information transmission is possible only if preferences are sufficiently aligned.

Chwe (1999), Persico (2004), Gerardi et al. (2005), and Chwe (2006) propose information eliciting by means of distorting the decision or manipulating agents' payoffs via a bet on the state of the world. In this chapter, we are not interested in optimal mechanisms. Instead, we want to study a widely used one: discussions followed by majority voting. We identify conditions for the existence of equilibria in which information is fully aggregated, i.e in which the decision is the same as if all signals were common knowledge.

The papers discussed above share the common finding that full information aggregation is impossible if preferences in the committee are too heterogeneous. One conclusion one might be tempted to draw from this result is that a small group of decision makers with aligned interests can make better decisions than a larger group of decision makers with conflicting interests. This conclusion would imply that a representative democracy might be an inferior decision-making institution. The present chapter suggests that this view on representative democracies is too pessimistic. We show that full information aggregation can be possible in spite of preference heterogeneity.

We follow the existing literature and model the decision procedure as a deliberating process followed by a voting phase, but we study a different information environment. We assume that the decision-relevant private information is verifiable. That is, committee members are assumed to be aware of facts which can be proven. This assumption corresponds well to committees whose members are experts on the issue. An example would be a board of examiners who have to propose a candidate for a grant. Decision-relevant facts could for instance be the candidate's performance in individual examiners' courses which is verifiable, but private information.[1] We provide more examples for our setting in the next section. Transmission of verifiable information has been studied e.g. in Milgrom and Roberts (1986) and Banerjee and Somanathan (2001), where an informed party or several informed parties try to influence a decision maker by revealing information. In contrast to these papers, informed parties participate in the decision process. Moreover, their preferred decision may also depend on the other agents' information. One could presume that players are able to force each other to reveal verifiable information. In our model, there is no means for players to affect each other's payoffs except for the decision they collectively make.

In the basic model, we assume that preferences are common knowledge. This assumption is well in accordance with situations in which decision makers are elected in order to represent different interests (like in a representative democracy), or with situations in which members are sent into the committee as a representative of an affected group (like in hiring committees). Our analysis provides a more optimistic view on the information aggregation potential of heterogeneous committees than previous work does. We show that an equilibrium exists in which information is perfectly aggregated. This is not the case in the games studied by Doraszelski et al. (2003) and Austen-Smith and Feddersen (2002). The reason is that in our set-up, committee members may be able to perfectly deduce the information of a voter who does not reveal it voluntarily. By not communicating his private information, a voter reveals that he possesses information he does not want to reveal. This contains exactly the same information as revealing the information itself. If information is soft (see also Austen-Smith (1990a,b)) the option to report false information destroys the opportunity to credibly transmit this information if it is indeed the truth.

Private information concerning preferences and soft information concerning the quality of the decision are good assumptions for novel and rare decision situations, whereas the approach in our basic model is well in accordance with

[1] It is likely that in reality, there is soft information on top of that. In this paper, we deal only with the aggregation of hard information. As soft information communication games always have babbling equilibria, we could argue that if there was soft information which somebody tried to communicate, nobody would listen.

committees consisting of experts who know each other, or whose interests can be inferred from their role in the committee. Examples are representative governments, hiring committees, or boards of directors.

We extent our basic model into two directions and derive conditions for full information aggregation in each extended set-up. First, we allow for the possibility that some agents are not endowed with decision-relevant information. Moreover, we examine an environment in which preferences are private information. Last, we combine these two modifications and consider a framework in which preferences are private information and in which there is the possibility that agents are not endowed with information.

In the modified versions of the model, full information aggregation is possible only if the preference parameter range is restricted. If committee members are not endowed with information with certainty, full information revelation in the communication stage is possible if and only if interests are completely aligned. The reason is that committee members with information which is unfavorable for their favorite decision can pretend to have no information. Strongly biased committee members prefer to conceal such information. However, we can show that there exists an equilibrium in which every committee member reveals at least one type of information, and each type of information is revealed by more than half of the committee members (if they possess such information) for cases in which the probability of receiving information is high enough.

If preferences are private information, there exists an equilibrium in which information is completely revealed if preference diversity is not too extreme. Committee members are uncertain about the majority's preferred alternative. As it is possible that the majority's interests are aligned with their own, committee members have an incentive to provide information. This is supported by a belief system with the property that information revelation does not harm in cases in which the majority's preferred decision deviates from one's own. Uncertainty about the majority's preferences may provide incentives for information revelation if committee members do not possess information with certainty. There exist equilibria with full information revelation for preference parameter constellations for which this is not the case if preferences are common knowledge.

The chapter is organized as follows. In the next section the model setup is presented. We derive the full information aggregation result of the basic model in Section 3.3. In Section 3.4 the extensions to the basic model are analyzed and conditions are derived under which full information aggregation is possible. In the final section we conclude and outline possible directions for future research.

3.2 The basic model

A collective decision $x \in \{a, b\}$ is made by majority voting without abstentions in a committee consisting of n members. For the ease of exposition (to avoid ties), let n be an odd number. Utility from the decision is state-dependent. There are two possible states of nature $\omega \in \{A, B\}$, and uncertainty about its realization. Ex ante both states are equally likely.

Each agent i receives a signal $\sigma_i \in \{\alpha, \beta\}$ which is correlated with the true state of the world:

$$prob\,\{\sigma_i = \alpha\,|\omega = A\} = prob\,\{\sigma_i = \beta\,|\omega = B\} = q,\ \tfrac{1}{2} < q < 1\ \forall i.$$

The signals are drawn independently conditional on the state. A signal contains verifiable information. Prior to voting, there is the possibility to communicate within the committee. Verifiability of information implies that committee members cannot invent information: they can either report the information they are endowed with or stay silent.

Examples for this decision environment are the following:

- $x \in$ {stick to the status quo, implement a reform}; $\omega \in$ {the reform causes higher costs than benefits, the benefits outweigh the costs}; $\sigma_i \in$ {presumptive evidence for either state: a certain group looses surely but little, the reform worked in a neighbor state, etc.}.

- $x \in$ {hire a new researcher, not}; $\omega \in$ {researcher is brilliant, researcher is mediocre}; $\sigma_i \in$ {presumptive evidence for either state: researcher has a joint paper in a leading field journal; researcher performed badly at a conference, etc.}.

- $x \in$ {conviction of a defendant; acquittal}; $\omega \in$ {defendant is guilty, defendant is innocent}; $\sigma_i \in$ {presumptive evidence for either state: defendant would have had a good reason to commit the crime, defendant has never been conspicuous so far, etc.}.

The timing is as follows:

1. Nature draws the state of the world and an imperfect signal for every agent.
2. Agents may reveal their signal to the other agents.
3. Agents vote. The alternative which receives the most votes is implemented.

The solution concept is Perfect Bayesian Nash equilibrium. That is, at each possible node of the game in which a player is asked to take an action, the action is required to be a best response to the other players' strategies given the beliefs, and beliefs shall be consistent with equilibrium strategies.

3.2.1 Agents

Agents derive state-dependent utility from the collective decision, $U_i = u_i(x, \omega)$. They are Bayesians and seek to maximize expected utility taking into account all available information. Let $p_i(\omega = A)$ denote the probability which agent i assigns to state of the world A given the information available to him.
Agent i's expected utility from decision a is:

$$p_i(\omega = A)\, u_i(a, A) + (1 - p_i(\omega = A)) u_i(a, B),$$

and from b :

$$(1 - p_i(\omega = A))\, u_i(b, B) + p_i(\omega = A)\, u_i(b, A).$$

Throughout the analysis, we assume a certain degree of homogeneity in preferences, which ensures that the desirability of decision a weakly increases in the probability that state A realized for each agent.[2]

Assumption 3.1 $u_i(a, A) + u_i(b, B) - u_i(a, B) - u_i(b, A) > 0\ \forall i.$

Agent i prefers the implementation of a over b, iff

$$p_i(\omega = A) > \frac{u_i(b, B) - u_i(a, B)}{u_i(a, A) + u_i(b, B) - u_i(a, B) - u_i(b, A)}. \tag{3.1}$$

Denote $\theta_i = \frac{u_i(b,B) - u_i(a,B)}{u_i(a,A) + u_i(b,B) - u_i(a,B) - u_i(b,A)}$. Arrange the names of agents $i = 1, \ldots n$ such that $\theta_1 \leq \theta_2 \leq \cdots \leq \theta_n$, and denote the median type, $\theta_{\frac{n+1}{2}}$, with

[2] If Assumption 3.1 does not hold, the number of voters who prefer decision a over decision b may not be monotone in the probability that state A has realized. The analysis then requires to consider all possible shapes which this relationship may have.

θ_m. Following (among others) Feddersen and Pesendorfer (1998), θ_i is called i's threshold of doubt. Agent i prefers a over b if and only if he assesses the probability that the state of the world is A higher than his threshold of doubt. Agents $i : \theta_i < 0$ prefer decision a in both states of the world, and agents $j : \theta_j > 1$ prefer decision b in both states of the world. The present paper allows for a larger preference parameter range than most of the existing literature (e.g. Fedderson and Pesendorfer (1998) or Doraszelski et al. (2003)), where attention is restricted to the case $\theta_i \in [0, 1]$.

Agents $l : \theta_l \in [0, 1]$ agree on the decision that should be made under certainty. Hence there are incentives to pool private information in order to get a better estimate about the true state of the world. However, heterogeneous thresholds of doubt potentially cause disagreement at the time the decision has to be taken. Therefore, agents may not want to reveal their information if this could cause the politically decisive voter to vote against their preferred alternative.

Preferences are common knowledge. We preclude the implementation of transfer schemes. Reasons for this restriction are that either (i) there exists no authority which is able to collect the transfers after the decision was implemented and the state of the world was learned, (ii) the state of the world is not verifiable, and/or (iii) individual rationality and budget constraints cannot be satisfied simultaneously.[3]

3.2.2 Information processing

As utility is state-dependent, agents would like to condition their choice on the state of the world. The state of the world is not observable, but correlated with individual signals. Agents use the information about the realization of the signals for updating their beliefs concerning the realization of the state of the world. Firstly, each agent observes a signal, which alters his beliefs about the state and about the distribution of signals held by the other committee members. Secondly, agents observe the communication outcome and therewith the realization of a subset of the signals. Moreover, they are able to interpret the actions of those committee members who did not reveal their information. Last, each agent

[3] If we could align preferences over the collective decision by a payoff manipulation $u_i(x, \omega)$ via a transfer scheme $\{t_i(x, \omega), x = a, b; \omega = A, B; i = 1, \ldots n\}$, beneficial information aggregation would be no problem.

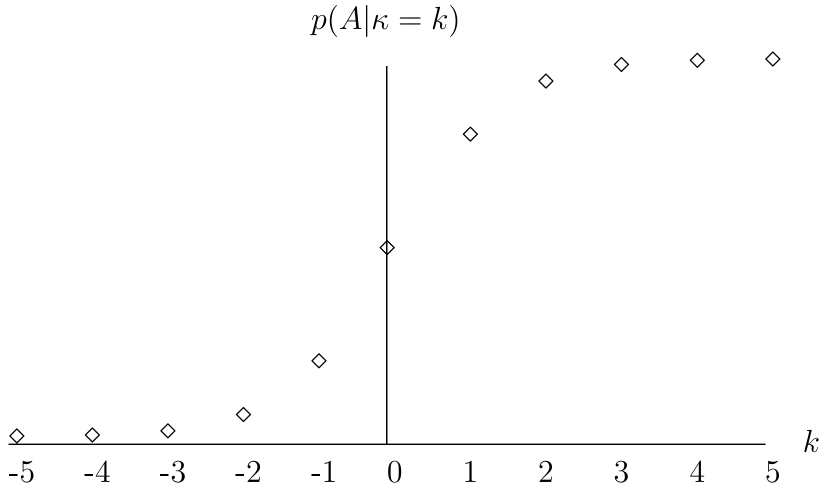

Figure 3.1: Probability that the state of the world is A given the evidence k.

is aware of the fact that in equilibrium his vote affects the outcome only for particular realizations of the other voters' signals.

Suppose for the moment that the realization of private information $\sigma_i,\ i = 1, \ldots n$ is common knowledge. Since the information environment is symmetric (i.e. $p(A) = 1/2$, $prob(\sigma_i = \alpha|\omega = A) = prob(\sigma_i = \beta|\omega = B)\ \forall i$), the only information which matters for updating beliefs with respect to the realization of the state of the world is the difference between the number of α-signals and β-signals.[4] Denote this random variable with κ. We have $\kappa = \sum_{i=1}^{n} 1_{\sigma_i=\alpha} - 1_{\sigma_i=\beta}$, where $1_{\sigma_i=\hat{\sigma}} = 1$ if $\sigma_i = \hat{\sigma}$ and 0 else. Bayesian updating yields:

$$p\left(A\,|\kappa = k\right) = \frac{q^k}{q^k + (1-q)^k}. \tag{3.2}$$

Figure 3.1 depicts the probability that the state of the world is A given there are k more α-signals than β-signals.[5]

If the realization of some $\sigma_j,\ j = 1, \ldots n$, (and hence of κ) is not known, agent i has to form beliefs $\mu_i(\sigma_j),\ i \neq j$, with respect to these realizations, incorporating all available information. We denote with κ_{-i} the difference between the number of α-signals and β-signals held by committee members except for i. The beliefs

[4]See Appendix A.

[5]Note that (as n is an odd number) κ assumes even values with positive probability only if some agents do not receive information, see Sections 3.4.1 and 3.4.3.

held with respect to the other agents' signals can be transformed into a belief regarding κ_{-i}.[6]

$$\mu_i(\kappa_{-i} = k) = \begin{cases} \sum_{J \subseteq \{1,...n\} \backslash \{i\}: |J| = \frac{n-1+k}{2}} \prod_{j \in J} \mu_i(\sigma_j = \alpha) \prod_{l \in \{1,...n\} \backslash \{J \cup i\}} \mu_i(\sigma_l = \beta), \\ \qquad \text{for } k \in E\{-n+1, ..., n-1\} \\ 0, \qquad \text{for } k \notin \{-n+1, ..., n-1\}, \end{cases} \tag{3.3}$$

where $E\{x, ..., y\}$ denotes the set of even integers between (including) x and y. Taking his own signal and $\kappa = \kappa_{-i} + 1_{\sigma_i = \alpha} - 1_{\sigma_i = \beta}$ into account, i's belief regarding κ is given by:

$$\mu_i(\kappa = k) = \begin{cases} \mu_i(\kappa_{-i} = k+1), & \text{if } \sigma_i = \alpha \\ \mu_i(\kappa_{-i} = k-1), & \text{if } \sigma_i = \beta. \end{cases} \tag{3.4}$$

Agent i's belief regarding the state of the world is then given by:

$$p_i(\omega = A) = \sum_{k \in O\{-n,...,n\}} \mu_i(\kappa = k) p\left(A \,|\kappa = k\right), \tag{3.5}$$

where $O\{-n, ...n\}$ is the set of odd integers between $-n$ and n, and $\mu_i(\kappa)$ is updated whenever the agent receives new information. There are several stages at which the beliefs can be updated.[7] Firstly, the agent learns his own signal σ_i. Secondly, the other agents' observed communication actions contain information. Moreover, agents may be able to deduce information through equilibrium voting strategies. They care only about their vote when it is pivotal, hence they update their beliefs using the information contained in this event. If the event that i is pivotal never occurs in equilibrium, basically any beliefs can be assigned, provided that they do not contradict Bayes' Law.

Figure 3.2 depicts a possible path of belief-updating for an agent in a committee with five agents. The illustrations represent agent i's beliefs with respect to κ at different stages of the game. Note that positive probability is assigned only to odd values for κ because there is an odd number of signals. Ex ante, the probability that k_α agents receive α-signals and $k_\beta = n - k_\alpha$ agents receive β-signals is given by $\binom{n}{k_\alpha}(\frac{1}{2}q^{k_\alpha}(1-q)^{k_\beta} + \frac{1}{2}q^{k_\beta}(1-q)^{k_\alpha})$. That is, prior to the receipt of information, probability is assigned to $\kappa = k$ according to:

$$\binom{n}{\frac{n+k}{2}}(\tfrac{1}{2}q^{\frac{n+k}{2}}(1-q)^{\frac{n-k}{2}} + \tfrac{1}{2}q^{\frac{n-k}{2}}(1-q)^{\frac{n+k}{2}}),$$

[6] See Appendix A.

[7] For convenience, we use the term 'updating' even if the updating does not change an agents' assessment.

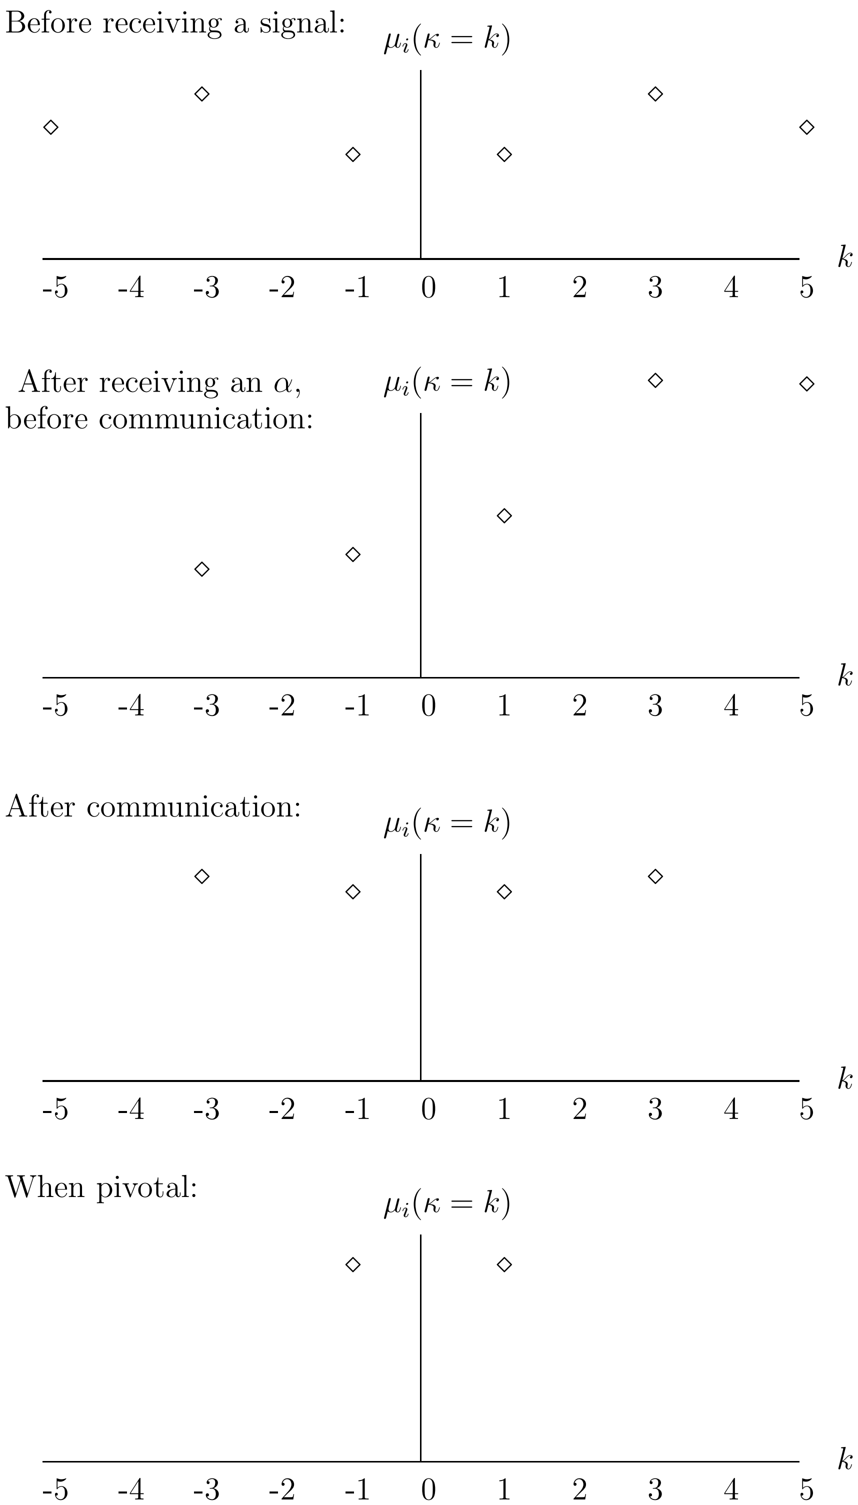

Figure 3.2: Possible path of updating in the course of the game with 5 players.

where $k \in O\{-n, \dots n\}$.

After having learned that his signal is α, i assigns positive probability only to $\kappa \in O\{-n+2, \dots n\}$, and assigns probability to $\kappa = k$ according to:

$$\frac{\binom{n-1}{\frac{n+k}{2}-1}\left(q^{\frac{n+k}{2}}(1-q)^{\frac{n-k}{2}} + q^{\frac{n-k}{2}}(1-q)^{\frac{n+k}{2}}\right)}{\sum_{k' \in O\{-n+2,\dots n\}} \binom{n-1}{\frac{n+k'}{2}-1}\left(q^{\frac{n+k'}{2}}(1-q)^{\frac{n-k'}{2}} + q^{\frac{n-k'}{2}}(1-q)^{\frac{n+k'}{2}}\right)}.$$

Suppose that i was shown one β in the communication stage, and that those who stayed silent planned to do so for both types of signals.[8] Then, i can exclude $\kappa = 5$ and conclude $p_i(\omega = A) = 1/2$. At the stage of voting, he again updates his beliefs, assigning positive probability only to those κ which – given the voting strategy profile – may render his vote decisive. Suppose that the voter who revealed the β-signal votes for b, that one of the remaining voters votes for a irrespective of his information, and that the other two voters vote informatively, that is each of them votes for a if his signal is α, and votes for b if his signal is β. Then, i's vote is decisive if and only if those who vote informatively have opposing signals.

3.2.3 The communication stage

The agents are allowed to reveal the signal they received from nature prior to the voting stage. As a signal contains verifiable information, agents cannot lie about their information. A communication strategy γ_i for an agent i is a plan whether to report his information (σ_i) or to remain silent (s) for each signal he may receive. Communication takes place simultaneously and is observed by all voters.[9]

Let C denote the set of possible outcomes of the communication stage. An outcome of the communication stage is denoted $c = (c_1, c_2, \dots c_n)$, where $c_i \in \{\alpha, \beta, s\}\ \forall i = 1, \dots n$. Denote with $\mathcal{A}(c)$ the set of agents who revealed α-signals, $\mathcal{A}(c) = \{i \in \{1, \dots n\} : c_i = \alpha\}$. Define $\mathcal{B}(c)$ and $\mathcal{S}(c)$ analogously. The most important summaries of the information provided in the communication stage are the number of revealed α-signals, $k_\alpha(c) = |\mathcal{A}(c)|$, the number of revealed

[8] If a voter j reveals an α-signal with a higher probability than a β-signal, beliefs formed after observing j's silence would take this into account, assigning higher probability to $\sigma_j = \beta$.

[9] The full information aggregation result (Proposition 3.1) does not hinge upon the assumption of simultaneous communication.

β-signals, $k_\beta(c) = |\mathcal{B}(c)|$, and the number of unrevealed signals, $k_s(c) = |\mathcal{S}(c)|$. Denote with $k(c) = k_\alpha(c) - k_\beta(c)$ the number of revealed α-signals in excess of the number of β-signals. We call $k(c)$ *evidence*, and say that the communication stage produced evidence for A if $k(c) > 0$, evidence for B if $k(c) < 0$ and no evidence if $k(c) = 0$.

Having observed j's communication action, i updates his belief regarding the realization of j's signal. Denote with $\mu_i(\sigma_j = \alpha|c_j = \hat{c})$ the probability which i assigns to $\sigma_j = \alpha$ given j's communication action $\hat{c}$. Because communication strategies are restricted to either truthful revelation or no revelation, we have that $\mu_i(\sigma_j = \alpha|c_j = \alpha) = 1$, and $\mu_i(\sigma_j = \alpha|c_j = \beta) = 0$. Beliefs regarding the signals of voters $j : j \in \mathcal{S}(c)$ must be consistent with these agents' communication strategies along the equilibrium path, and have to respect Bayes' Rule off the equilibrium path.

3.2.4 The voting stage

Agents vote simultaneously and without abstentions for an alternative to be implemented. The alternative which gets the most votes is implemented. Every agent takes into account all the information available to him, that is the own signal σ_i, the communication outcome c, and what he can learn through equilibrium play. In particular, agents base their votes on being pivotal.

A voting strategy v_i for agent i is a plan for which alternative to vote, for all possible outcomes of the communication stage and for each signal he may receive. Allowing for mixed voting strategies, we have $v_i : \{\alpha, \beta\} \times C \to [0, 1]$, where $v_i(.)$ denotes the probability to vote for a.

Definition 3.1 *A voting strategy which satisfies*

$$v_i\,(.) = \begin{cases} 1, & \textit{if } p_i\,(\omega = A) > \theta_i \\ 0, & \textit{if } p_i\,(\omega = A) < \theta_i \end{cases}$$

is called Bayesian sincere voting.

Recalling (3.1), it is easy to see that a Bayesian sincere voting strategy maximizes i's expected utility. Moreover, if the voting strategy profile allows the event that i is pivotal to occur, the beliefs $\mu_i(\kappa)$ are well-defined and Bayesian sincere voting is the only utility-maximizing strategy. However, if an agent is

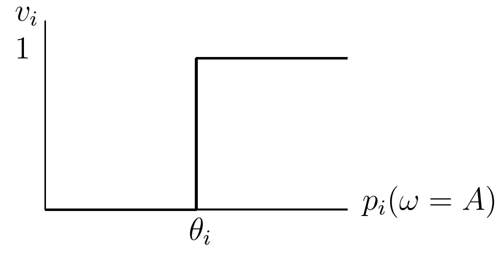

Figure 3.3: Bayesian sincere voting strategy

never pivotal, any voting action is utility-maximizing, since none has an effect. Therefore, Bayesian sincere voting is always in the best response set at the stage of voting for every player. Figure 3.3 depicts a Bayesian sincere voting strategy for voter i. Assume again that the realization of the signals – and hence the realization k of κ – is common knowledge. Then, using Bayesian sincere voting strategies, agents i with thresholds of doubt $\theta_i < p(A|\kappa = k)$ vote for a, and agents $j : \theta_i > p(A|\kappa = k)$ vote for b. An agent's threshold of doubt reflects how much evidence for state of the world A must be presented to the agent such that he supports alternative a. For convenience, we assume that the voter with the median preference type is never indifferent between a and b if the realization of κ is known with certainty. This assumption is made without loss of generality to avoid case differentiations. Alternatively, if he is indifferent between the two alternatives, we could restrict attention to equilibria in which the median preference type takes a particular action, say vote for a.

Assumption 3.2 *$\exists k_m \in \{-n, \dots n\}$ such that $p(A|\kappa = k_m) < \theta_m < p(A|\kappa = k_m + 1)$.*

3.3 Full information aggregation

Both, preferences over the set of alternatives and beliefs regarding the state of the world may differ among voters. Voter i with preferences θ_i does not want alternative a to be implemented as long as he assesses the probability that the state of the world is A to be smaller than θ_i. Communication allows for the possibility to influence the politically decisive voter – and consequently the collective decision – in one's own favor. Whether or not a voter has the possibility to influence the decision in his favor depends on the kind of information he is endowed with.

Thus, the mere possibility to communicate may cause information exchange, even if the committee members do not talk to each other.

It is easy to see that an equilibrium exists in which the information is fully revealed and taken into account by every voter when making the voting choice. The result is stated in the following proposition.

Proposition 3.1 *The following strategy profile is a Perfect Bayesian Nash equilibrium of the communication-and-voting-game: All agents reveal their information in the communication stage. All agents assign probability 1 to* $\kappa = k(c)$*. Agents whose thresholds of doubt are smaller than* $p(A|\kappa = k(c))$ *vote* a*, the other agents vote* b*. If* $c_i = s$ *is observed, and* $\theta_i \geq \theta_m$*, voters* $j \neq i$ *assign probability 1 to* $\sigma_i = \alpha$ *and play voting strategies as if* i *had revealed* α*. If* $c_i = s$ *is observed, and* $\theta_i < \theta_m$*, voters* $j \neq i$ *assign probability* 0 *to* $\sigma_i = \alpha$ *and play voting strategies as if* i *had revealed* β*.*

Proof. We have to distinguish two cases: (i) k_m is odd, and (ii) k_m is even. Consider agent $i : \theta_i \geq \theta_m$. The only deviation in the communication stage which has an effect is to conceal β. In case (i), this deviation will change the outcome only if $\kappa_{-i} = k_m + 1$. If i sticks to his equilibrium communication action, all voters believe (know) that $\kappa = k_m$, and assess the probability that the state of the world is A to be $p(A|\kappa = k_m)$. As $\theta_m > p(A|\kappa = k_m)$, a majority votes for b. If i conceals β, the other agents believe that $\kappa = k_m + 2$ and consequently that the probability that the state of the world is A is $p(A|\kappa = k_m + 2)$. $\theta_m < p(A|\kappa = k_m + 1)$, hence a majority votes for a. In case (ii), i's revelation affects the outcome only if $\kappa_{-i} = k_m$. If i sticks to his equilibrium communication action, all voters believe that $\kappa = k_m - 1$, and a majority votes for b. If i conceals β, the other agents believe that $\kappa = k_m + 1$ and consequently a majority votes for a. Since $\theta_i \geq \theta_m$, agent i prefers b over a in case his deviation has an effect. Thus, i has no incentive to deviate in the communication stage. The voting strategy is Bayesian sincere. Thus, there exists no profitable deviation. A similar argument applies to the case $\theta_i < \theta_m$. Note that beliefs are consistent with Bayes' Rule. Q.E.D.

The intuition for the result is as follows. Full information revelation is possible because committee members may apply the following reasoning: I know who you are, I know what you want, and I know that you know something. If you don't tell me what you know, then I suppose that you have information which is unfavorable for your favorite decision because otherwise you would have told me.

In equilibrium, every agent has an incentive to reveal at least one type of information, given he possesses it: Agents who are more biased towards a than the median preference type prefer alternative a whenever the majority prefers a. Hence, they have an incentive to reveal α-signals in order to convince the majority of alternative a. Agents who are biased towards b have an incentive to provide β-signals. By not revealing what he knows, an agent reveals that he does not know anything he wants to reveal. As all agents know something, staying silent reveals exactly the same information as the revelation of the information itself, and turns out to be an inconsequential action.

3.4 Extensions to the basic model

In this section, we modify the basic setting along two dimensions and identify conditions under which full information revelation in the communication stage is possible. First we allow for the possibility that agents do not possess decision-relevant information with a certain probability. In this case, agents with unfavorable information can pool with agents who have no information. Hence, it may not always be possible to perfectly deduce the information endowment of a committee member who does not talk. Next, we consider the case that preferences are private information. Then it is impossible to assume 'worst news' in case of a player's silence because what would be bad news is private information. In both modifications of the basic model, full information aggregation is possible only if committee members' preferences are sufficiently aligned. However, there may still be a large range of preference parameters for which full information aggregation is possible. Interestingly, if there is the possibility that agents do not receive information, then a full information aggregation equilibrium exists for a larger preference parameter range if preferences are private information then if they are common knowledge.

We identify conditions for the existence of equilibria in which all signals are revealed in the communication stage and call these equilibria *full information revelation equilibria*. One could presume that the existence of such an equilibrium is not necessary for full information aggregation, as non-revealed information may still be aggregated in the voting stage. We will show that if the extended model does not have a full information revelation equilibrium, then there neither exists a full information aggregation equilibrium.

3.4.1 Possibility of receiving no signal

In the basic model, full information aggregation is possible because committee members may apply the reasoning "I know what you want, and I know that you know something. If you don't tell me what you know, then I suppose that you have information which is unfavorable for your favorite decision." Here, we eliminate the "I know that you know something" part from this line of argumentation. We assume that each committee member receives a signal with probability $\delta < 1$. If a voter is not endowed with information, we denote $\sigma_i = \emptyset$. We assume that it is impossible to verify $\sigma_i = \emptyset$, and that this event is equally likely in both states of the world. That is, the message spaces remain the same as in the basic model for those agents who receive information, whereas those who do not receive information are restricted to staying silent.

We modify nature's moves in the following way:

1. Nature determines the set of agents $\Delta \subseteq \{1, \ldots n\}$ that she will endow with signals.
2. Nature draws the state of the world.
3. Nature draws a signal for each agent $i \in \Delta$.

Suppose that full information revelation as in Proposition 3.1 was part of an equilibrium in this game as well. Then, in this equilibrium, if a voter i does not reveal information in the communication stage, the only admissible belief for the other voters is to assign probability 1 to $\sigma_i = \emptyset$. In an equilibrium in which all the available information is revealed in the communication stage, $\mu_i(\kappa)$ assigns probability 1 to $\kappa = k(c)$ for all agents i. Hence, it is rational for voter i to vote for a if $p(A|\kappa = k(c)) > \theta_i$, and to vote for b if $p(A|\kappa = k(c)) < \theta_i$.

Consider agent i, $\sigma_i = \alpha$, and assume that agents $-i$ reveal any signal they receive from nature, and hold the belief that if $c_i = s$, then $\sigma_i = \emptyset$ with probability 1. Agent i is pivotal with signal α if and only if $\kappa_{-i} = k_m$. In this case the revelation would cause a majority to vote for a, while a majority votes for b given only the evidence κ_{-i}. If agents $-i$ expect i to reveal any information nature endows him with, they think he has no information if he stays silent, and assign probability 1 to $\kappa = k_m$ in case i deviates. A concealment would not be noticed. However, whether or not he reveals his own signal to the other agents, i would know that $\kappa = k_m + 1$, and assign probability $p(A|\kappa = k_m + 1)$ to state of the world A. Hence, he prefers the revelation of α if and only if $\theta_i < p(A|\kappa = k_m+1)$. With the same line of reasoning, we conclude that an agent $i : \sigma_i = \beta$ reveals his signal if and only if all the other agents reveal their signals if $\theta_i > p(A|\kappa = k_m)$.

Proposition 3.2 *Consider a communication-and-voting-game in which each voter is endowed with information with probability $\delta < 1$. There exists a full information revelation equilibrium if and only if $\theta_i \in [p(A|\kappa = k_m), p(A|\kappa = k_m + 1)]$ $\forall i$.*

Full information revelation is possible if and only if all committee members agree with the median preference type which decision should be made given the presented evidence, i.e. if there is essentially no preference heterogeneity. However, there may still be considerable information aggregation even in the presence of strongly diverging interests. Suppose that there is an agent $i : \theta_i < p(A|\kappa = k_m)$. Agent i disagrees with the majority only insofar as they implement b in some cases in which i would better like a. This is why i does not want to show a β-signal if he is endowed with it. However, agent i might be willing to reveal an α-signal. The following proposition states that the existence of an equilibrium with considerable information aggregation is guaranteed if δ is high enough.

Proposition 3.3 *Consider a communication-and-voting-game in which each voter is endowed with information with probability $\delta < 1$. There is a lower bound on δ, $\delta' < 1$, such that for all $\delta > \delta'$, there is a k with $|k| < |k_m|$ such that, an equilibrium exists in which at least $\frac{n+1}{2}$ agents reveal α-signals if they are endowed with them, at least $\frac{n+1}{2}$ agents reveal β-signals if they are endowed with them, and all agents vote for a if $k(c) \geq k+1$ and vote for b if $k(c) \leq k$.*

We present the proof of Proposition 3.3 in Appendix A. We show that given the beliefs induced by a certain strategy profile, it is possible to construct this strategy profile such that the strategies are mutually best responses and beliefs are indeed correct. For small values of δ, the construction is possible for some preference parameter constellations, but cannot be applied for the general case. The reason is that as δ goes to zero, beliefs conditional on providing decisive information converge to $p(\omega = A|\kappa = k)$ and $p(\omega = A|\kappa = k + 1)$, but are also sensitive to the communication strategy profile. Hence, we cannot guarantee that an equilibrium exists in which equilibrium beliefs concentrate around the median preference type. The equilibrium identified in Proposition 3.3 has the property that more than half of the agents reveal α-signals if they are endowed with them, and more than half of the agents reveal β-signals if they are endowed with them. The decision rule is such that less extreme evidence has to be presented in order to change the decision as is necessary for the median preference type to change his mind in the full information case. The reason is that the revelation of an additional α-signal affects beliefs in two ways, directly via the correlation, and indirectly because less unrevealed information remains.

As $\delta \to 1$, $k \to k_m/2$ for the class of equilibria identified in Proposition 3.3. Note that for $\delta = 1$, if half of the agents reveal either type of information, and the communication stage yields evidence $k(c)$, it follows that $\kappa = 2k(c)$. Hence, for $\delta = 1$, the decision will be a if and only if $\kappa > k_m$, which implies full information aggregation as in Proposition 3.1 also in this class of equilibria. Although we may have considerable information aggregation in the absence of a full information revelation equilibrium, full information aggregation is impossible for $\delta < 1$.

Proposition 3.4 *Consider a communication-and-voting-game in which each voter is endowed with information with probability $\delta < 1$. If a full information revelation equilibrium does not exist, then there exists no full information aggregation equilibrium.*

Proof. Full information aggregation requires the following. If $\kappa \leq k_m$, then there must be at least $\frac{n+1}{2}$ b-votes. If $\kappa > k_m$, then there must be at least $\frac{n+1}{2}$ a-votes. This in turn is feasible only (i) if a voter who did not reveal β votes b and would have voted a if $\sigma_i = \emptyset$, and (ii) a voter who did not reveal α votes a and would have voted b if $\sigma_i = \emptyset$. As beliefs must be consistent, an agent who votes b anticipates that in case his vote is pivotal, then $\kappa = k_m$, and a pivotal agent who votes a infers that $\kappa = k_m + 1$. Hence, the required voting strategies are consistent with equilibrium play only if $\theta_i \in [p(A|\kappa = k_m), p(A|\kappa = k_m + 1)]\ \forall i$. Then a full information revelation equilibrium exists. Q.E.D.

3.4.2 Private information concerning preferences

In this section, we assume that not only individual signals but also preferences are private information. Therewith, we eliminate the "I know what you want"-part from the line of argumentation underlying Proposition 3.1. Agent i's preference parameter θ_i is drawn according to a commonly known probability function $\phi(\theta_i)$, which is assumed to be identical for all $i = 1, \ldots n$. The realization of θ_i is i's private information. We denote the distribution function of individual types with $\Phi(\theta_i)$.

Nature's moves are now as follows:

1. Nature draws the agents' types.

2. Nature draws the state of the world.

3. Nature draws a signal for each committee member.

We are interested in conditions for the existence of a full information revelation equilibrium. Clearly, agents with types above 1 (respectively below 0) would never reveal information which makes the choice of a (respectively b) more likely. If all the other agents reveal their information and vote Bayesian sincerely, the revelation of an α-signal (respectively the revelation of a β-signal) necessarily has this effect (if any). Hence, in order that a full information revelation equilibrium exists, the support of $\phi(\theta_i)$ must be bounded. That is, there must exist $\theta_{min} > 0$, $\theta_{max} < 1$ such that $\theta_i \in [\theta_{min}, \theta_{max}]\,\forall i$. In particular, for every i, there must be an integer $k_{\theta_i} \in E\{-n+1, \ldots, n-1\}$ such that $\theta_i \in [p(A|\kappa = k_{\theta_i} - 1), p(A|\kappa = k_{\theta_i} + 1)]$, i.e. such that agent i prefers decision a for all $\kappa_{-i} > k_{\theta_i}$ and prefers decision b for all $\kappa_{-i} < k_{\theta_i}$. For $\kappa_{-i} = k_{\theta_i}$, he prefers a if his own signal is α, and b if his own signal is β.

In the following, we assume the existence of a full information revelation equilibrium in which all agents vote Bayesian sincerely (which implies full information aggregation). If all the information is revealed during the communication stage, then $p_i(\omega = A) = p(A|\kappa = k(c))\,\forall i$, and hence the agents' Bayesian sincere voting strategies are unique. In case of remaining uncertainty about decision-relevant information, Bayesian sincere voting strategies are determined by the belief system μ. Hence, the incentive to reveal information - and therewith full information revelation in equilibrium - hinges upon the beliefs agents hold in case of a deviation (i.e. a concealment of information). To derive conditions for the existence of a full information revelation equilibrium, we specify beliefs μ which best support the full information revelation equilibrium. It suffices to specify these beliefs for the case that a single agent conceals his information in the communication stage.

Given that preferences are private information, it does not make sense to condition the beliefs in case of i's silence on i's name. However, beliefs can be conditioned on the communication outcome $k(c)$. Denote with $\mu_{-i}(k(c)) =$

κ_{-i}	revelation of α	has an effect iff $\theta_m \in$	benefits/ harms i iff
$< k_{\theta min}$	no effect	-	-
$k_{\theta min}$	$b \to a$	$[p(A\|\mu(k_{\theta min}))$, $p(A\|\kappa = k_{\theta min}+1)]$	$k_{\theta_i} < k_{\theta min}+1$ $k_{\theta_i} > k_{\theta min}+1$
$\hat{k}$	$b \to a$	$[p(A\|\mu(\hat{k}))$, $p(A\|\kappa = \hat{k}+1)]$	$k_{\theta_i} < \hat{k}+1$ $k_{\theta_i} > \hat{k}+1$
$k_{\theta max}$	$b \to a$	$[p(A\|\kappa = \mu(k_{\theta max}))$, $p(A\|\kappa = k_{\theta max}+1)]$	$k_{\theta_i} < k_{\theta max}+1$ $k_{\theta_i} > k_{\theta max}+1$
$> k_{max}$	no effect	-	-
κ_{-i}	revelation of β	has an effect iff $\theta_m \in$	benefits/ harms i iff
$< k_{\theta min}$	no effect	-	-
$k_{\theta min}$	$a \to b$	$[p(A\|k_{\theta min}-1)$, $p(A\|\mu(k_{\theta min}))]$	$k_{\theta_i} > k_{\theta min}-1$ $k_{\theta_i} < k_{\theta min}-1$
$\hat{k}$	$a \to b$	$[p(A\|\kappa = \hat{k}-1)$, $p(A\|\mu(\hat{k}))]$	$k_{\theta_i} > \hat{k}-1$ $k_{\theta_i} < \hat{k}-1$
$k_{\theta max}$	$a \to b$	$[p(A\|\kappa = k_{\theta max}-1)$, $p(A\|\mu(k_{\theta max}))]$	$k_{\theta_i} > k_{\theta max}-1$ $k_{\theta_i} < k_{\theta max}-1$
$> k_{max}$	no effect	-	-

Figure 3.4: Possible effects of i's revelation given information revelation by agents $-i$.

$\mu_{-i}(\kappa = k(c)+1|\kappa_{-i} = k(c), c_i = s)$ the out-of-equilibrium-belief agents $j \neq i$ assign to $\sigma_i = \alpha$ in case of i's silence given the communication outcome $k(c)$.

To see how these beliefs best support a full information revelation equilibrium, consider the possible effects of i's revelation of the two types of signals given agents $-i$ reveal their signals, and given beliefs $\mu_{-i}(k(c))$. These are illustrated in Figure 3.4. In this figure beliefs with respect to the state of the world given the communication outcome c and information revelation by all agents but i are (with slight abuse of notation) denoted with $p(A|\mu(k(c)) = (1-\mu_{-i}(k(c)))p(A|\kappa = k(c)-1) + \mu_{-i}(k(c))p(A|\kappa = k(c)+1)$. Consider agent $i : \theta_i = \theta_{min}$. Whenever this agent's revelation of an α-signal has an effect, this effect is beneficial for agent i. The reason is that (as i is most biased towards a) whenever the majority prefers a over b given all the available information, then i likes a better than b as well. However, the revelation of a β-signal can have a beneficial effect for i only if $\kappa_{-i} = k_{\theta_{min}}$. For $\kappa_{-i} < k_{\theta_{min}}$, all agents agree that b is the best choice, regardless of i's signal. For all $\kappa_{-i} > k_{\theta_{min}}$ the revelation of a β-signal can only harm i. Given the realization of the other agents' signals, $\kappa_{-i} = \hat{k}$, the revelation of a β-signal will change the majority decision from a to b if and only if the median of the preference types θ_m realized within the interval $[p(A|\kappa = \hat{k}-1), (1-\mu_{-i}(\hat{k}))p(A|\kappa = \hat{k}-1) + \mu_{-i}(\hat{k})p(A|\kappa = \hat{k}+1)]$, that is if the Bayesian sincere voting strategies for the majority prescribe to vote for b in case i reveals a β-signal, and prescribe to vote for a given the beliefs $\mu_{-i}(k(c))$ if agent i conceals his information. Note that the probability that the median voter's preference type realizes in the relevant range is highest, and hence the incentive to reveal a β-signal is strongest for agent $i : \theta_i = \theta_{min}$ if out-of-equilibrium-beliefs assign probability 1 to $\kappa = k_{min} + 1$ if $\kappa_{-i} = k_{min}$. Note also that given this belief, i's revelation of an α-signal has no effect on expected utility for $\kappa_{-i} = k_{min}$.

We now quantify the effect of information revelation versus information concealment in a full information revelation equilibrium on i's expected utility. First, we fix the realizations of the random variables and suppose that the revelation of a β-signal changes the majority decision from a to b, given $\kappa_{-i} = \hat{k}$. The effect on i's expected utility is:

$$\begin{aligned} & (1-p(A|\kappa = \hat{k}-1))(u_i(b,B) - u_i(a,B)) \\ & -p(A|\kappa = \hat{k}-1)(u_i(a,A) - u_i(b,A)) \\ = \; & (u_i(b,B) - u_i(a,B) + u_i(a,A) - u_i(b,A))(\theta_i - p(A|\kappa = \hat{k}-1)), \quad (3.6) \end{aligned}$$

which is positive for $(\theta_i > p(A|\kappa = \hat{k}-1))$ and proportional to $(\theta_i - p(A|\kappa = \hat{k}-1))$. This implies that whenever the smallest preference type gains from the revelation of a β-signal, then every other type gains as well. Symmetrically, whenever the highest preference type gains from the revelation of an α-signal, then every other type does so.

The probability which agent i assigns to the effect (3.6) on his expected utility (i.e. the probability assigned to the joint events $\kappa = \hat{k} - 1$ and $\theta_m \in [p(A|\kappa = \hat{k}-1), (1-\mu_{-i}(\hat{k}))p(A|\kappa = \hat{k}-1) + \mu_{-i}(\hat{k})p(A|\kappa = \hat{k}+1)]$) depends on his own private information, σ_i and θ_i.

Let $\phi_i^m(\theta')$ denote the probability which agent i assigns to the event that the median voter has a type θ' given his own preference type θ_i. It depends on his own type θ_i because i is part of the sample drawn from $\phi(\theta)$. $\phi_i^m(\theta')$ is given by (A.7)–(A.9) which are stated in Appendix A and depicted in Figure 3.5 for $n = 5$, and a uniform distribution on $[0.1, 0.9]$ of individual preference types. The figure

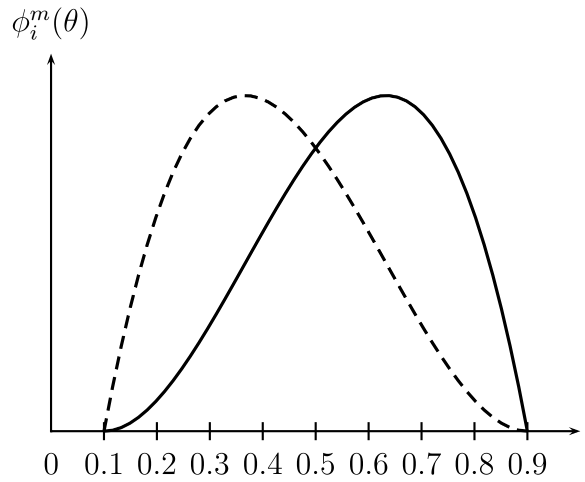

Figure 3.5: $\phi_i^m(\theta)$, for $n = 5$, $\phi(\theta) = U[0.1, 0.9]$, $\theta_i = 0.9$ (solid), and $\theta_i = 0.1$ (dashed).

shows $\phi_i^m(\theta)$ for $\theta_i = 0.1$ and $\theta_i = 0.9$.

Let $\mu_i(\kappa_{-i} = k|\sigma_i)$ denote the probability which agent i assigns to the event that the realization of the other agents' signals yield $\kappa_{-i} = k$ given his own signal σ_i. Note that an agent with an α-signal assigns a higher probability to high realizations of κ_{-i} than an agent with a β-signal because individual signals are correlated via the state of the world. $\mu_i(\kappa_{-i} = k|\sigma_i)$ is given by (A.10) and (A.11) in Appendix A and depicted in Figure 3.6 for $n = 5$ and $q = 0.8$. The figure shows $\mu_i(\kappa_{-i} = k|\sigma_i)$ for $\sigma_i = \alpha$ and $\sigma_i = \beta$.

Given that all agents $-i$ reveal their information and hold beliefs $\mu(k(c))$, the effect of i's revelation of a β-signal on expected utility is proportional to:

$$\sum_{k=k_{\theta_{min}}}^{k_{\theta_{max}}} \mu_i(\kappa_{-i} = k|\beta)(\theta_i - p(A|\kappa = k-1)) \int_{p(A|\kappa=k-1)}^{(1-\mu_{-i}(k))p(A|\kappa=k-1)+\mu_{-i}(k)p(A|\kappa=k+1)} \phi_i^m(\theta)d\theta, \quad (3.7)$$

and the effect of the revelation of an α-signal on i's expected utility is proportional to:

$$\sum_{k=k_{\theta_{min}}}^{k_{\theta_{max}}} \mu_i(\kappa_{-i} = k|\alpha)(p(A|\kappa = k+1) - \theta_i) \int_{(1-\mu_{-i}(k))p(A|\kappa=k-1)+\mu_{-i}(k)p(A|\kappa=k+1)}^{p(A|\kappa=k+1)} \phi_i^m(\theta)d\theta. \quad (3.8)$$

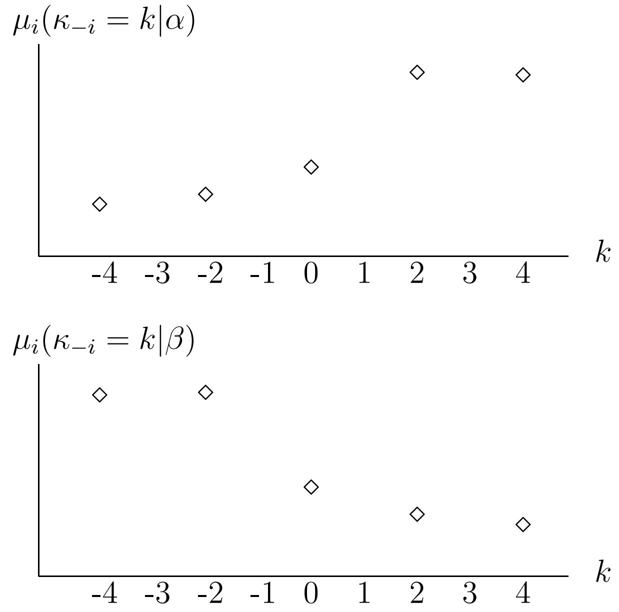

Figure 3.6: Beliefs regarding the other agents' signals given the own signal ($n = 5$, $q = 0.8$).

As argued above, (3.8) is unambiguously positive for $\theta_i = \theta_{min}$. Similarly, (3.7) is positive for $\theta_i = \theta_{max}$. In the following we derive conditions under which (3.7) is non-negative for $\theta_i = \theta_{min}$, and (3.8) is non-negative for $\theta_i = \theta_{max}$.
(3.7) is non-negative for $\theta_i = \theta_{min}$ iff

$$\mu_i(\kappa_{-i} = k_{\theta_{min}}|\beta)(\theta_{min} - p(A|\kappa = k_{\theta_{min}} - 1)) \int_{p(A|\kappa=k_{\theta_{min}}-1)}^{(1-\mu_{-i}(k_{\theta_{min}}))p(A|\kappa=k_{min}-1)+\mu_{-i}(k_{\theta_{min}})p(A|\kappa=k_{\theta_{min}}+1)} \phi^m_{min}(\theta)d\theta$$
$$\geq \sum_{k=k_{\theta_{min}}+2}^{k_{\theta_{max}}} \mu_i(\kappa_{-i} = k|\beta)(p(A|\kappa = k-1) - \theta_{min}) \int_{p(A|\kappa=k)}^{(1-\mu_{-i}(k))p(A|\kappa=k-1)+\mu(k)p(A|\kappa=k+1)} \phi^m_{min}(\theta)d\theta, \quad (3.9)$$

and (3.8) is non-negative for $\theta_i = \theta_{max}$ iff

$$\mu_i(\kappa_{-i} = k_{\theta_{max}}|\alpha)(p(A|\kappa = k_{\theta_{max}} + 1) - \theta_{max}) \int_{(1-\mu_{-i}(k_{\theta_{max}}))p(A|\kappa=k_{\theta_{max}}-1)+\mu_{-i}(k_{\theta_{max}})p(A|\kappa=k_{\theta_{max}}+1)}^{p(A|\kappa=k_{\theta_{max}}+1)} \phi^m_{max}(\theta)d\theta$$
$$\geq \sum_{k=k_{\theta_{min}}}^{k_{\theta_{max}}-2} \mu_i(\kappa_{-i} = k|\alpha)(\theta_{max} - p(A|\kappa = k)) \int_{(1-\mu_{-i}(k))p(A|\kappa=k-1)+\mu_{-i}(k)p(A|\kappa=k+1)}^{p(A|\kappa=k+1)} \phi^m_{max}(\theta)d\theta, \quad (3.10)$$

where $\phi^m_{min}(\theta)$ is given by (A.9), $\phi^m_{max}(\theta)$ is given by (A.7), $\mu_i(\kappa_{-i} = k|\beta)$ is given by (A.11), and $\mu_i(\kappa_{-i} = k|\alpha)$ is given by (A.10).

Obviously, (3.9) and (3.10) are necessary conditions for the existence of a full information revelation equilibrium. Note that in case agent i's expected utility is higher if he reveals β (α) than if he conceals the information, then the same is true for agent $j : \theta_j > \theta_i$ ($\theta_j < \theta_i$) because j benefits relatively more and looses relatively less than agent i whenever the revelation has an effect. Hence, conditions (3.9) and (3.10) are also sufficient for the existence of a full information revelation equilibrium. That is, for the existence of a full information revelation equilibrium in which agents vote Bayesian sincerely, it suffices to make sure that the type who is most biased towards alternative a is willing to reveal a β-signal and that the type who is most biased towards alternative b is willing to reveal an α-signal.

If $k_{\theta_{min}} = k_{\theta_{max}}$, then (3.9) and (3.10) hold for any $\mu_{-i}(k(c))$, as the left-hand-sides are zero. It should be intuitively clear that there is no incentive to conceal information in this case because $k_{\theta_{min}} = k_{\theta_{max}}$ implies that there is essentially no preference heterogeneity, i.e. voters agree on the mapping of information into the decision.

If $k_{\theta_{min}} < k_{\theta_{max}}$, then the right-hand-side of (3.9) increases, and the left-hand-side of (3.10) decreases in $\mu_{-i}(k_{min})$. Hence, out-of-equilibrium beliefs $\mu_{-i}(k(c))$ best support the full information revelation equilibrium if $\mu_{-i}(k_{\theta_{min}}) = 1$. Similarly, out-of-equilibrium-beliefs best support the full information revelation equilibrium if $\mu_{-i}(k_{\theta_{max}}) = 0$.

This observation yields sufficient conditions for the existence of a full information revelation equilibrium, stated in Proposition 3.5.

Proposition 3.5 *Consider a communication-and-voting-game in which the preference parameters θ_i are private information.*
(i) There exists a full information revelation equilibrium if the preference parameters are drawn from $[p(A|\kappa = k-1), p(A|\kappa = k+3)]$ for all agents i and some even integer $k \in \{-n+1, \ldots, n-3\}$.
(ii) There exists a full information revelation equilibrium if the preference parameters are drawn from $[p(A|k_{\theta_{min}}-1), p(A|k_{\theta_{min}}+1)] \cup [p(A|k_{\theta_{max}}-1), p(A|k_{\theta_{max}}+1)]$ for all agents i and some even integers $k_{\theta_{min}}, k_{\theta_{max}} \in \{-n+1, \ldots, n-1\}$.

Hence, full information aggregation is possible in heterogeneous committees if preference heterogeneity is not too severe. For $q = 0.8$, a sufficient condition

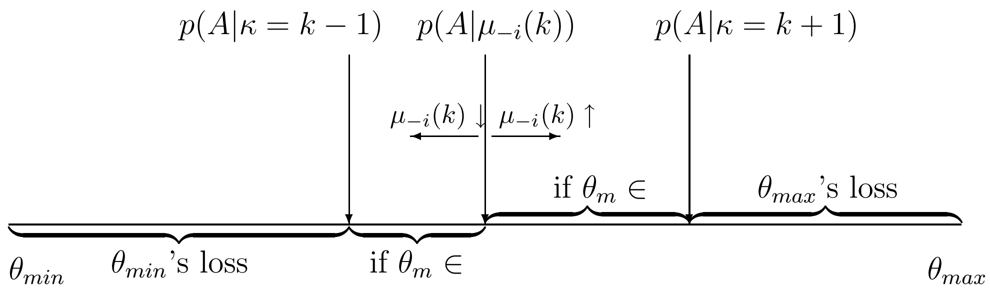

Figure 3.7: $\mu_{-i}(k)$ determines the probabilities with which i looses from signal revelation.

for the existence of a full information revelation equilibrium is that preference types are drawn from the interval $[0.2, 0.985]$. That is, the committee may have members who need to be at least 80% sure that the state of the world is B in order to support alternative b as well as agents who need to be 98.5% sure that the state of the world is A in order to support decision a (and any preference type in between). Moreover, full information aggregation is possible regardless of the quality of the signal if preference types assume only two values between 0 and 1, provided that there exist realizations of the signals for each type which convinces him of either alternative (this can be achieved by increasing the number of committee members).

If the conditions stated in Proposition 3.5 hold, out-of-equilibrium-beliefs can be defined such that a revelation has an effect only if the preferences of the median voter are aligned with the own preferences. However, even if information revelation has unfavorable effects in some cases, the effect may still be favorable in expectation, such that (3.9) and (3.10) hold for more heterogeneous preferences (as measured by $\theta_{max} - \theta_{min}$) than in part (i) of the proposition and more general preference distributions than in part (ii). It is difficult to identify the out-of-equilibrium-beliefs which best support a full information revelation equilibrium. Figure 3.7 illustrates how the out-of-equilibrium-beliefs affect the incentives for information revelation for the most biased committee members for $\kappa_{-i} = k \neq k_{\theta_{min}}, k_{\theta_{max}}$. By decreasing $\mu_{-i}(k)$ for some $k \neq k_{min}, k_{max}$, the incentive to reveal a β-signal increases for type θ_{min} because the probability that he will incur the loss $p(A|\kappa = k-1) - \theta_{min}$ decreases. At the same time, the incentive for type θ_{max} to reveal an α-signal decreases because he will incur the loss $\theta_{max} - p(A|\kappa = k+1)$ with a higher probability. The smaller (higher) k, the higher the loss incurred by type θ_{max} (θ_{min}) in case the revelation of a β-signal (α-signal) has an effect. From this point of view, $\mu_{-i}(k)$ should be high for small realizations of κ_{-i}, and low for high values. However, type θ_{min}, endowed with a β-signal, assigns less probability to both events than type θ_{max} with an α-signal, (i) the realization of a high κ_{-i}, and (ii) the realization of the median type in the relevant range (see Figures 3.5 and 3.6). Which of the two forces is stronger depends on the parameters of the model. As the necessary conditions for the existence of a full information revelation equilibrium can only be derived with knowledge of the most favorable out-of-equilibrium-beliefs, they cannot be stated without further specification of the model. We provide an example in Appendix A, where we show that the sufficient conditions for the existence of a full information revelation equilibrium stated in Proposition 3.5 are not necessary.

We can again exclude other full information aggregation equilibria, if the full information revelation equilibrium does not exist.

Proposition 3.6 *Consider a communication-and-voting-game in which the preference parameters θ_i are private information. If a full information revelation equilibrium does not exist, then there exists no full information aggregation equilibrium.*

Proof. A necessary condition for a strategy profile to be a full information aggregation equilibrium is that the decision is responsive to each signal. This requires informative voting by those who did not reveal their signals. Then, all a-voters draw the same inferences out of being pivotal, and all b-voters draw the same inferences when they are pivotal. Hence, there is a k such that $p(A|\kappa = k-1) \leq \theta_i \leq p(A|\kappa = k+1)\ \forall i$. A full information revelation equilibrium exists. Q.E.D.

If preferences are private information, the preference parameter range for which a full information revelation equilibrium exists is smaller than in the common knowledge case (where it is unbounded), but larger than in the case in which committee members are informed about other members' preferences but can pretend to have no information. The reason is that the median voter may agree even with the most biased committee members. In the next section we show that this effect may allow for a larger preference parameter range in case committee members can pretend to have no information.

3.4.3 Private preferences, possibility of receiving no signal

If there is the possibility of receiving no signal, we cannot support the full information revelation equilibrium with out-of-equilibrium-beliefs anymore, because no action can be identified as being out-of-equilibrium. As pointed out in Section 3.4.1, in a full information revelation equilibrium the only admissible belief regarding i's signal when observing $c_i = s$ is $\sigma_i = \emptyset$. The necessary and sufficient conditions for the existence of a full information revelation equilibrium are – analogously to (3.9) and (3.10):

$$\mu_i(\kappa_{-i} = k_{min} + 1|\beta)(\theta_{min} - p(A|\kappa = k_{min})) \int_{p(A|\kappa=k_{min})}^{p(A|\kappa=k_{min}+1)} \phi^m_{min}(\theta)d\theta$$
$$\geq \sum_{k=k_{min}+2}^{k_{max}+1} \mu_i(\kappa_{-i} = k|\beta)(p(A|\kappa = k-1) - \theta_{min}) \int_{p(A|\kappa=k-1)}^{p(A|\kappa=k)} \phi^m_{min}(\theta)d\theta,$$

$$\mu_i(\kappa_{-i} = k_{max}|\alpha)(p(A|\kappa = k_{max} + 1) - \theta_{max}) \int_{p(A|\kappa=k_{max})}^{p(A|\kappa=k_{max}+1)} \phi^m_{max}(\theta)d\theta$$
$$\geq \sum_{k=k_{min}}^{k_{max}-1} \mu_i(\kappa_{-i} = k|\alpha)(\theta_{max} - p(A|\kappa = k+1)) \int_{p(A|\kappa=k)}^{p(A|\kappa=k+1)} \phi^m_{max}(\theta)d\theta,$$

and $\theta_{min} \in [p(A|k_{min}), p(A|k_{min} + 1)]$ and $\theta_{max} \in [p(A|k_{max}), p(A|k_{max} + 1)]$.

Note that the necessary condition for the existence of a full information revelation equilibrium in case preferences are common knowledge, $k_{min} = k_{max}$, (see Proposition 3.2) is sufficient here. As in the previous section, committee members are uncertain about the median voter's preferences. If the probability that the majority has the same interests as one's own is high enough, committee members have an incentive to reveal their information even if there are (potentially) conflicts of interest. We illustrate this possibility by means of an example.

Consider a committee with three members, and suppose each of them receives a signal from nature with probability δ. If an agent receives a signal, the signal is correct with probability 0.8. Preference parameters θ_i are drawn from $[0.2 + \epsilon, 0.8 - \epsilon]$ according to a uniform distribution. Suppose agents 1 and 2 reveal their information. If both reveal an α-signal, then all agents agree the decision should be a regardless of agent 3's information. If both reveal a β-signal, then all agents agree the decision should be b regardless of agent 3's information. Suppose agent 3 is endowed with a β-signal. This information will be pivotal if (i) $\theta_m \in [0.2, 0.5]$ and (a) agents 1 and 2 reveal different signals, or (b) none of the other agents received information, or (ii) $\theta_m \in [0.5, 0.8]$ and one of the other agents revealed an α-signal and the other agent did not receive information. If $\theta_3 \in [0.2, 0.5]$, the revelation of the β-signal is beneficial in case (i), but not in case (ii). If $\theta_3 \in [0.5, 0.8]$, the revelation is beneficial in both cases. Agent $3 : \theta_3 = 0.2 + \epsilon$ (and hence all other types) has an incentive to reveal a β-signal if

$$\begin{aligned} \epsilon \cdot \tfrac{3}{4}(\delta^2 \cdot 2 \cdot 0.8 \cdot 0.2 + (1-\delta)^2) &> (0.3 - \epsilon) \cdot \tfrac{1}{4} \cdot 2\delta(1-\delta)2 \cdot 0.8 \cdot 0.2 \\ \Leftrightarrow \epsilon &> \frac{\delta(1-\delta)}{5\delta^2 + \frac{125}{8}(1-\delta)^2 + \frac{10}{3}\delta(1-\delta)}. \end{aligned}$$

Because of the symmetry of the parameter constellation, the condition for the revelation of an α-signal for preference type $0.8 - \epsilon$ is the same.

For values of δ close to 0 or 1, a full information revelation equilibrium exists for very small ϵ. Independently of δ, the existence of a full information revelation equilibrium in this example is guaranteed if $\theta_{min} = 0.25$, and $\theta_{max} = 0.75$. Note that for this preference parameter range, there exists no full information revelation equilibrium in case preferences are common knowledge. Summing up, we have the following proposition.

Proposition 3.7 *Consider two communication-and-voting-games Γ^c, Γ^p in which each voter is endowed with a signal with probability $\delta < 1$. The games are identical except that in Γ^c, preferences are common knowledge, and in Γ^p, preferences are private information.*
(i) If Γ^c has a full information revelation equilibrium, then Γ^p has a full information revelation equilibrium.
(ii) There are parameter constellations such that Γ^p has a full information revelation equilibrium, whereas no full information revelation equilibrium exists for Γ^c.

For $\delta < 1$ the existence of a full information revelation equilibrium hinges upon voluntary information revelation by the players. As we have shown, the incentive to do so may be greater in the light of uncertainty about the majority's preferences. For $\delta = 1$, the event that a committee member does not reveal information arises in a full information revelation equilibrium only off the equilibrium path. Information revelation can be supported by out-of-equilibrium beliefs. If preferences are common knowledge, out-of-equilibrium beliefs can be conditioned on the preferences of the deviating player, whereas in the private information case, this is impossible.

3.5 Conclusion

This paper provides a first step towards the analysis of committees who deal with verifiable information, and whose members have conflicting interests. We identified conditions under which all decision-relevant information is revealed at the communication stage and taken into account at the stage of voting. If preferences are common knowledge, and every committee member is endowed with information with certainty, then there exists an equilibrium with these properties independently of the extent of preference heterogeneity.

If preferences are private information, then there exists a full information revelation equilibrium if preference heterogeneity is not too severe. If preferences are common knowledge, but agents are endowed with information only with a certain probability, then full information aggregation is possible if and only if all voters agree how the information should be mapped into a decision, i.e. if there are no conflicts of interest. Moreover, if there is the possibility of receiving no signal then preferences being private information allows for a larger extent of preference heterogeneity than the common knowledge case. The reason is the possibility that the majority can have the same interests as oneself.

Our results may be used to analyze the quality of collective decisions in several extended frameworks. First of all, it is worth studying incentives for information acquisition. In the present paper, information comes for free and in the basic model every committee member possesses private information with certainty. The impossibility to lie about the realization of private signals allows the committee to deduce a member's information even if this member does not want the committee to be aware of it. This could weaken incentives to acquire information in the first place beyond the usually found free riding problem.

If preferences are homogeneous 'enough', we can expect efficient information aggregation. However, information aggregation may be a problem if preferences are too heterogeneous (if there is the possibility that some agents are not endowed with information and/or preferences are private information). Agents might want to exclude agents who have preferences which are too distinct from their own from communication, while sharing their information with more like-minded. There may be demand for a device which allows agents to match into a homogeneous subgroup in order to pool information more efficiently. It would be interesting to compare the efficiency of a system in which information is pooled within subgroups (which may be interpreted as political parties) who are represented by a single voice in the voting stage to the efficiency of information aggregation within a direct democracy.

We used the simple majority rule as the decision mechanism. For homogeneous preferences, there exists a unique best decision rule (Costinot and Kartik, 2006). An interesting extension would be to take a mechanism design perspective in a setting with heterogeneous preferences. Suppose for instance an alternative needs a fraction $q > 1/2$ of the votes in order to be implemented. If no alternative gets this fraction, then the status quo is maintained. Then, information α is pivotal in two cases: for changing the decision from status quo to a, and for changing it from b to status quo. Full information revelation might be possible for parameter ranges for which it is impossible using simple majority rule. The optimal mechanism must trade-off the provision of incentives to reveal information versus the risk of maintaining the status quo to often.

In our model, individual signals are – conditional on the state of the world – independent random variables. This is a good assumption if committee members have different areas of expertise. In other cases, it might be more appropriate to allow for the possibility that the information contained in the agents' signals partially overlaps. An example would be the hiring committee, where some of the candidates' characteristics are more easily observable than others. The information environment could be modeled as a set of verifiable signals, containing information about the alternatives at hand, out of which nature draws a subset for each committee member. In such a setting (again referring to the hiring committee), it would be particularly interesting to allow for a manipulation of nature's moves (influenced by the candidates' actions) and to study the interaction with committee members' information acquisition efforts.

Chapter 4

Communication in Committees: Who should listen?

4.1 Introduction

When a firm decides whether to enter a new market or not, it usually faces uncertainty about costs and benefits of market entry. If costs and benefits accrue asymmetrically to decision makers and may be observed privately, an information aggregation problem arises. If the marketing division reaps the largest part of the benefits from market success, and the other decision makers (e.g. the production unit) bear most of the costs, the marketing division might be reluctant to inform the other decision makers about high costs. How should the information flow within the group of decision makers be organized in order to make use of private information, and to enter the market when it is worth entering and to stay out when it is not?

In this paper, we compare two alternative decision procedures. One of them allows decision makers to make public speeches before they decide via majority voting. In the other one, communication takes place within groups of decision makers with the same interests (for instance the marketing division) before the representatives of the groups make a decision via majority vote. We call the first decision procedure the *open debate mechanism*, and the latter the *group debate mechanism*. In a political economy context, we can think of the former as a parliament, and the latter as a party system. We are interested in the decision quality induced by the two mechanisms, measured in terms of expected social surplus.

The main difference between the two decision procedures is that members of one group have no access to the information which is revealed within the other group in the group debate, whereas they do have access to this information in the open debate. Although restricting access to communication seems to be a waste of information given that it is accessible, it might not have been accessible if everybody had listened. If some decision makers strongly favor one alternative, then an individual who is less biased towards that alternative may be cautious to reveal information which makes their votes for the alternative even more likely. Moreover, as committee members anticipate that information in favor of their preferred alternative will be concealed, they presume that an agent who does not talk probably has information which confirms their predisposition. As a consequence, choices are less sensitive to information in the open debate mechanism than in the group debate mechanism.

Our framework can be applied to collective decision frameworks in which the consequences of the decision are borne asymmetrically by the decision makers. Consider for instance decisions concerning public good provision, say a municipality's decision whether to build a dyke or not. If the dyke is financed by taxes, all decision makers are affected symmetrically by the costs. However, those who live near to the coast enjoy larger benefits if the dyke is built than those living in the inland. Another example is a faculty's decision whether to hire a researcher with a focus on applied econometrics or on economic theory. The costs of the hiring decision will be borne symmetrically by the faculty members. However, benefits may accrue asymmetrically to the decision makers, depending on their own research interests. Conflicting interests among decision makers may hinder the information exchange within the decision-making institution, and the efficient aggregation of decision-relevant information.

There is a growing literature on information aggregation in committees (for a survey, see Chapter 2). If communication among committee members is possible prior to making a choice, the committee is able to make (weakly) better decisions than without the possibility to communicate. In fact, a committee of homogeneous members can make efficient use of the available information (see e.g. Coughlan (2000)). If preferences are heterogeneous, the possibility to communicate still (weakly) increases decision quality (Doraszelski et al. (2003)). One may presume that the more communication possibilities there are, the better the decision quality. In this chapter, we show that this intuition may be misleading. If committee members have heterogeneous preferences, decision quality may be higher when the decision makers exchange information within a subgroup of members with aligned interests rather than if they talk to the entire committee.

Previous work has emphasized that information sharing during debate is problematic if committee members have conflicting interests (see e.g. Austen-Smith (1990a,b), Doraszelski et al. (2003), and Meirowitz (2005)). Piketty (1999) provides a survey on the role of political institutions for information aggregation in the political process. Austen-Smith and Feddersen (2006) study the impact of the voting rule on the extent of information sharing in debate. In Ottaviani and Sorensen (2001), the optimal sequence of statements is derived in a setting in which decision makers want to appear well-informed. So far, the question to whom the debate should be addressed has not received attention. In the present chapter, we explore this question.

There are several contributions which study the question with whom to communicate prior to making a choice, when the informed agent does not take part in the decision process (e.g. Dewatripont and Tirole (1999), Dur and Swank (2005), and Gerardi et al. (2005)). Wolinsky (2002) allows for communication among experts with identical interests which differ from those of the decision maker. He shows that it may be beneficial for the decision maker to allow only partial communication among the experts. In contrast to these papers, we study a decision framework in which decision makers themselves may be (partially) informed about the state of the world, and have to decide whether to share their information with other decision makers.

Maug and Yilmaz (2002) study the performance of a two-class voting mechanism compared to voting within the entire group. They also find that a separation of the committee into homogeneous subgroups may increase decision quality. Maug and Yilmaz (2002) do not allow for communication, hence votes must reflect private information in order to make use of it. A two-class voting mechanism performs better in case of preference heterogeneity than voting within the whole committee because more voters base their decisions on information. The reason is that an individual vote is pivotal only if a majority of the other group votes in favor. Therewith, equilibrium information can be transmitted via the majority rules to make voters, who otherwise do not respond to their private information, more responsive. In the model presented in this chapter, communication is possible. Here, too much information is transmitted through equilibrium strategies in the open debate in the sense that a committee member's silence is interpreted as a signal against his predisposition. By precluding the interpretation of communication actions (by preventing their observation), individual votes are more responsive to information in the group debate mechanism.

The chapter is organized as follows. In the next section we present a parsimonious model of the decision environment within which the two mechanisms are studied. In Sections 4.3 and 4.4 the equilibria and the induced decision quality in the group debate mechanism and in the open debate mechanism are derived. We compare the two mechanisms and state our main result in Section 4.5. We conclude in Section 4.6.

4.2 The model

A group of four individuals has to decide whether to implement a project (enter a market, buy a public good, reform the welfare state) or not. (Per capita) costs (c) and benefits (b) arising from implementation are uncertain, and accrue asymmetrically to the committee members. We assume that there are two types of committee members. Committee members 1 and 2 are "low types" whose preferences can be represented by the following utility function:

$$U_l = \begin{cases} (1-\alpha)b - c, & \text{if the reform is implemented} \\ 0, & \text{else.} \end{cases}$$

Committee members 3 and 4 are "high types" whose preferences can be represented by the following utility function:

$$U_h = \begin{cases} (1+\alpha)b - c, & \text{if the reform is implemented} \\ 0, & \text{else.} \end{cases}$$

We assume that $\alpha > 0$. Hence, high types reap higher benefits from the project than low types do, and costs are shared equally. This fits for instance the dyke-building decision, where people who live near to the coast benefit more than people in the inland. In the market entry example, the marketing division might benefit more from market entry than the production division. The parameter α will be our measure of preference heterogeneity.

Costs and benefits can be either high or low, giving rise to four possible states of the world, $(b^l, c^l), (b^l, c^h), (b^h, c^l), (b^h, c^h)$, which are equally likely ex ante. Let $b^l = c^l = 1$ and $b^h = c^h = h > 1$. To make it a decision problem with only partially conflicting interests, we assume that $\alpha < \frac{h-1}{h}$. All players agree to implement the project in state (b^h, c^l), and not to implement in state (b^l, c^h). Low types prefer implementation only in state (b^h, c^l), and high types prefer implementation in all states except for (b^l, c^h). The stochastic structure implies that ex ante, high types are in favor of implementation, and low types are against implementation.

Prior to entering the decision procedure, an agent may receive a signal from nature, which contains perfect information about either the costs or the benefits, and can be credibly transmitted to other agents (depending on the decision procedure). Let δ be the probability to receive a signal, which is the same for all agents, and assume $0 < \delta < 1$. If an agent receives a signal, he will be informed about the costs or the benefits with equal probability. Individuals can infer information from the other players' actions during the game. Let $\mu_i(c^l)$ denote the probability agent i assigns to $c = c^l$, and let $\mu_i(b^h)$ be be the probability agent i assigns to $b = b^h$. Figure 4.1 summarizes the agents' expected payoffs from implementation for the cases of imperfect information about the state of the world.

We consider two decision procedures, the open debate mechanism, and the group debate mechanism. The open debate mechanism works as follows. In a communication stage, each committee member has the opportunity to reveal his signal (if endowed with one) to the other committee members. After the communication stage, each committee member casts a vote for or against implementation. Majority wins. In case of a tie, a fair coin is tossed. The group debate mechanism works as follows. In a communication stage, the low types have the opportunity to reveal their signals to each other, and the high types have the opportunity to reveal their signals to each other. After the communication stage, a representative of each group (say agent 1 for the low type group and agent 3 for the high type group) casts a vote for or against implementation. Again, majority wins, and in case of a tie, a fair coin is tossed. The two mechanisms are illustrated in Figure 4.2. Information revelation is observed by members inside the same box, but not outside the box.

In these games, a strategy for a player i consists of a revelation strategy γ_i, i.e. a plan which prescribes which signals to reveal if endowed with them, and – if i is supposed to cast a vote – a voting strategy v_i, i.e. a plan which vote to cast, A low type i's conditional expected utility from implementation:

Observed info	Conditional expected payoff	positive iff
b^l	$-\alpha - (1-\mu_i(c^l))(h-1)$	never
b^h	$(1-\alpha)h - \mu_i(c^l) - (1-\mu_i(c^l))h$	$\mu_i(c^l) > \frac{\alpha h}{h-1}$
c^l	$(1-\alpha)\big(\mu_i(b^h)h + (1-\mu_i(b^h))\big) - 1$	$\mu_i(b^h) > \frac{\alpha}{(1-\alpha)(h-1)}$
c^h	$(1-\alpha)\big(\mu_i(b^h)h + (1-\mu_i(b^h))\big) - h$	never
$\emptyset$	$(h-1)(\mu_i(b^h)(1-\alpha)h - 1 + \mu_i(c^l)) - \alpha$	$\mu_i(c^l) > \frac{h-1-\alpha}{h-1} - \mu_i(b^h)(1-\alpha)$

A high type i's conditional expected utility from implementation:

Observed info	Conditional expected payoff	positive iff
b^l	$\alpha - (1-\mu_i(c^l))(h-1)$	$\mu_i(c^l) > \frac{h-1-\alpha}{h-1}$
b^h	$(1+\alpha)h - p(c^l) - (1-\mu_i(c^l))h$	always
c^l	$(1+\alpha)\big(\mu_i(b^h)h + (1-\mu_i(b^h))\big) - 1$	always
c^h	$(1-\alpha)\big(\mu_i(b^h)h + (1-\mu_i(b^h))\big) - h$	$\mu_i(b^h) > \frac{h-1-\alpha}{(h-1)(1+\alpha)}$
$\emptyset$	$(h-1)(\mu_i(b^h)(1+\alpha)h - 1 + p(c^l)) - \alpha$	$\mu_i(c^l) > \frac{h-1+\alpha}{h-1} - \mu_i(b^h)(1+\alpha)$

Figure 4.1: Conditional expected payoffs from implementation.

Open debate mechanism: Group debate mechanism:

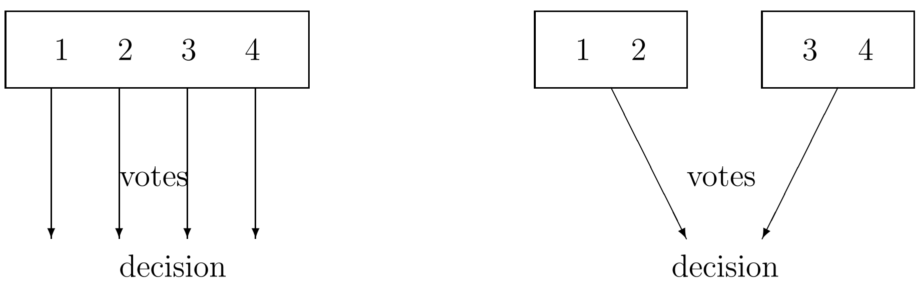

Figure 4.2: Open debate and group debate mechanism.

for each signal he may get and each communication outcome he might observe. Denote with γ a revelation profile, and with v a voting profile.

Our solution concept is Perfect Bayesian Nash equilibrium. That is, at each possible node of the game in which a player is asked to take an action, the action is required to be a best response to the other players' strategies given the beliefs. Beliefs shall be consistent with equilibrium strategies and Bayes' Rule on the equilibrium path, and shall not violate Bayes' Rule off the equilibrium path.

We measure decision quality in terms of expected social surplus, which is defined as the expected sum of payoffs. Expected social surplus depends only on the probability of implementation in state (b^h, c^l) and the probability of rejection in state (b^l, c^h). In the other states, committee members disagree, and the social surplus is zero for both decisions.

4.3 The group debate mechanism

In the voting stage, each of the two representatives casts a vote for or against implementation. A vote for (against) implementation increases (decreases) the probability of implementation by $\frac{1}{2}$ independently of the other representative's vote. Hence, a representative cannot infer any information from casting a pivotal vote. He votes for implementation if and only if the conditional expected payoff from implementation is positive, given the beliefs about the state of the world which were generated in the communication stage.

Group members have the same preferences. Given that the representative's vote maximizes his expected payoff, there is no incentive for a group member to conceal information. There may be equilibria in which some information is not revealed. Suppose for instance that α is so high that the low type representative votes for implementation only if he is sure that the state of the world is (b^h, c^l). Then, there is no need to report b^l because the representative will vote against implementation anyway. In that sense, the information is decision-irrelevant. However, there is no equilibrium in which decision-relevant information is not reported to the representative. Because a signal can only be revealed by a player who possesses that signal, the only admissible belief for the representative when observing an out-of-equilibrium revelation is to believe it. Thus, for any putative equilibrium in which decision-relevant information is concealed, revealing the information is a profitable deviation.

Without loss of generality, we can restrict attention to the equilibrium in which information is fully revealed within the groups. For communication outcomes which do not allow the identification of the state of the world, expected payoffs from implementation are given in Figure 1, where the beliefs $\mu_i(.)$ coincide with the priors for all communication outcomes. The equilibria of the group debate mechanism are characterized in the following lemma.

Lemma 4.1 *Consider the group debate mechanism. Any equilibrium has the following properties.*
The representative of the low type group votes for implementation (i) if he has observed b^h and c^l, and (ii) if he has observed c^l, and $\alpha \leq \frac{h-1}{h+1}$, (iii) if he has observed b^h, and $\alpha \leq \frac{h-1}{2h}$. Otherwise, he votes against implementation.
The representative of the high type group votes against implementation (i) if he has observed b^l and c^h, (ii) if he has observed b^l, and $\alpha \leq \frac{h-1}{2}$, and (iii) if he has observed c^h, and $\alpha \leq \frac{(h-1)}{h+1}$. Otherwise, he votes in favor of implementation.

We are interested in the probability of implementation if the state of the world is (b^h, c^l). Given that benefits are high and costs are low, whatever the representative of the high type group might learn during the communication stage, he will vote in favor of the project. Hence, if the representative of the low type group votes in favor of implementation, the probability of implementation is 1. Otherwise, it is $\frac{1}{2}$. If $\alpha \leq \frac{h-1}{2h}$, the low type representative votes in favor of implementation if and only if at least one of the agents in the low type group receives a signal. If $\frac{h-1}{2h} < \alpha \leq \frac{h-1}{h+1}$, the low type representative votes in favor of implementation if and only if at least one of them receives the information that costs are low. In case $\alpha > \frac{h-1}{h+1}$, the representative votes for implementation if and only if the group receives information about both, costs and benefits. The following lemma quantifies the probability of implementation in state (b^h, c^l) for the three cases.

Lemma 4.2 *Consider the group debate mechanism, and suppose that benefits are high and costs are low.*
(i) If $\alpha \leq \frac{h-1}{2h}$, the probability of implementation is

$$1 - \tfrac{1}{2}(1-\delta)^2.$$

(ii) If $\frac{(h-1)}{2h} < \alpha \leq \frac{(h-1)}{h+1}$, the probability of implementation is

$$1 - \tfrac{1}{2}\left(1 - \tfrac{\delta}{2}\right)^2.$$

(iii) If $\alpha > \frac{h-1}{h+1}$*, the probability of implementation is*

$$\tfrac{1}{2} + \left(\tfrac{\delta}{2}\right)^2.$$

Now, consider the case that benefits are low and costs are high. The representative of the low type group votes against implementation for any possible communication outcome. For $\alpha \leq \frac{h-1}{h+1}$, the representative of the high type group votes against implementation if and only if the group receives at least one signal. If $\frac{h-1}{h+1} < \alpha \leq \frac{h-1}{2}$, he votes against implementation if and only if at least one of the group members receives the information that benefits are low. If $\alpha > \frac{h-1}{2}$, he votes against implementation if and only if the group learns the state of the world. Lemma 4.3 quantifies the probability of rejection for the three cases.

Lemma 4.3 *Consider the group debate mechanism, and suppose that benefits are low and costs are high.*
(i) If $\alpha \leq \frac{h-1}{h+1}$*, the probability of rejection is*

$$1 - \tfrac{1}{2}(1-\delta)^2.$$

(ii) If $\frac{h-1}{h+1} < \alpha \leq \frac{h-1}{2}$*, the probability of rejection is*

$$1 - \tfrac{1}{2}\left(1 - \tfrac{\delta}{2}\right)^2.$$

(iii) If $\alpha > \frac{h-1}{2}$*, the probability of rejection is*

$$\tfrac{1}{2} + \left(\tfrac{\delta}{2}\right)^2.$$

From Lemmata 4.2 and 4.3, we can infer decision quality for the group debate mechanism, stated in the following proposition.

Proposition 4.1 *Expected social surplus generated in the group debate mechanism is:*

$$\begin{array}{ll}
(h-1)(1-(1-\delta)^2), & \textit{if } \alpha \leq \frac{h-1}{2h} \\
(h-1)(1-\frac{1}{2}\left(\left(1-\frac{\delta}{2}\right)^2 + (1-\delta)^2\right), & \textit{if } \frac{(h-1)}{2h} < \alpha \leq \frac{(h-1)}{h+1} \\
(h-1)\frac{1}{2}\left(\delta + \left(\frac{\delta}{2}\right)^2\right), & \textit{if } \frac{h-1}{h+1} < \alpha \leq \frac{h-1}{2} \\
(h-1)2\left(\frac{\delta}{2}\right)^2, & \textit{if } \alpha > \frac{h-1}{2}.
\end{array}$$

4.4 The open debate mechanism

In this section, we analyze the strategic interaction in the open debate mechanism. There is a class of equilibria in which all players vote for (against) implementation regardless of the communication outcome. Voting strategies are mutually best responses, because no single vote has an effect on the decision. It follows that any communication strategy is part of a best response. In this class of equilibria, the probability of implementation is 1 (0) in any state of the world. The group debate mechanism obviously yields a lower (higher) probability of implementation in state (b^h, c^l), and a higher (lower) probability of rejection in state (b^l, c^h). Because all states are equally likely ex ante, expected social surplus is higher in the group debate mechanism.

For a more convincing statement concerning the superiority of the group debate mechanism, however, we should compare the best equilibrium outcome of the open debate mechanism to the equilibrium outcome in the group debate mechanism. Hence, in the following we restrict attention to equilibria in which each voter votes for the alternative which maximizes his expected utility conditional on the information available to him. The term "equilibrium" will refer only to those strategy profiles which satisfy this additional criterion. Given that all players take into account all the available information at the stage of voting, decision quality will be the higher, the more information is revealed in the communication stage. Hence, we further restrict attention to equilibria that are *most revealing*.

Definition 4.1 *An equilibrium (γ^*, v^*) is most revealing if there is no equilibrium (γ', v') in which all players reveal all signals which they reveal in γ^*, and at least one player reveals a signal which he conceals in γ^*.*

With open debate, information revelation is observed by all players. Hence, a player's statement does not only affect the behavior of the like-minded player, but also that of the players who have different interests. Therefore, players may prefer to conceal information. In equilibrium, each player i's beliefs at the stage of voting $\mu_i(c^l), \mu_i(b^h)$ take into account equilibrium revelation strategies as well as equilibrium voting strategies. Note that whenever i finds himself in an information set in which any player has revealed c^l, whether on or off the equilibrium path, the only admissible belief for i yields $\mu_i(c^l) = 1$, because c^l can be revealed only by a player who has observed c^l, and c^l can only be observed if $c = c^l$.

Analogously, we have $\mu_i(c^l) = 0$ in an information set in which c^h was revealed, $\mu_i(b^h) = 1$ if b^h was revealed, and $\mu_i(b^h) = 0$ if b^l was revealed.

Restricting attention to most revealing equilibria, we presume that a player i conceals information only if it is a profitable deviation from a putative equilibrium in which players $-i$ play the same revelation strategies, and i is assumed to reveal the information. In order to identify such a profitable deviation in the communication stage, we have to study the effect of the deviation on the other individuals' actions in the voting stage. If a player i conceals information in the communication stage, this has an effect on i's expected payoff only if there is at least one voter j, a signal for this voter j, and announcements by players $-i$ such that if i reveals his information, j chooses a different voting action than if i stays silent. Obviously – as all players cast expected payoff maximizing votes in equilibrium – this can only be the case if the announcements by players $-i$ do not yield perfect information about the state of the world. Consider the players' expected payoffs for communication outcomes in which a concealment of information can possibly have an effect (see Figure 4.1).

A low type j's vote can be affected by concealing information only if an information set is reached in which (i) j has observed only b^h, (ii) j has observed only c^l, or (iii) j has observed nothing at all. Denote these information sets with $I_l(b^h)$, $I_l(c^l)$, and $I_l(\emptyset)$, respectively. Similarly, a high type j's vote can be affected by concealing information only if an information set is reached in which (a) j has observed only b^l, (b) j has observed only c^h, or (c) j has observed nothing at all. Denote these information sets with $I_h(b^l)$, $I_h(c^h)$, and $I_h(\emptyset)$, respectively. It is intuitive that no player has an incentive to conceal information in favor of his preferred alternative. This is stated in the following lemma. A proof of the lemma can be found in Appendix B.

Lemma 4.4 *Consider the open debate mechanism. In a most revealing equilibrium, low types reveal b^l and c^h, and high types reveal b^h and c^l.*

High types conceal at most b^l and c^h in a most revealing equilibrium, and low types conceal at most b^h and c^l. As high types have the same preferences, a high type's intention when concealing information can only be to trigger a vote for implementation by a low type, and a low type's intention can only be to trigger a vote against implementation by a high type. An immediate consequence of Lemma 4.4 is that these effects on a player's vote cannot be achieved if the player is uninformed at the stage of voting. This is stated in the following lemma. A proof can be found in Appendix B.

Lemma 4.5 *In a most revealing equilibrium of the open debate mechanism, if a low (high) type has not observed any information – neither privately nor during the communication stage – he votes against (in favor of) implementation.*

Concealing information has an effect only if it has a sufficient impact on a player's beliefs at the stage of voting. A high type's non-disclosure of c^h can have a desired effect on a low type j's vote only if j reaches an information set $I_l(b^h)$, and $\mu_j(c^l) > \frac{\alpha h}{h-1}$ at this information set. However, the non-disclosure can also affect the other high type k's vote, namely if he reaches an information set $I_h(b^l)$ and $\mu_k(c^l) > \frac{h-1-\alpha}{h-1}$ at this information set. Similarly, a low type's non-disclosure of c^l can have a desired effect on a low type j's vote only if j reaches an information set $I_h(b^l)$, and $\mu_j(c^l) < \frac{h-1-\alpha}{h-1}$ at this information set. The non-disclosure affects the other low type k's vote, if k reaches an information set $I_l(b^h)$ and $\mu_k(c^l) < \frac{\alpha h}{h-1}$ at this information set.

If no player conceals any information about the costs in a putative equilibrium, $\mu_i(c^l) = \frac{1}{2}$ $\forall i$ in any information set in which i is not perfectly informed about the costs. There is no incentive for any player to deviate from the putative equilibrium by concealing information about the costs, if $\frac{1}{2} < \frac{\alpha h}{h-1}$, and $\frac{1}{2} > \frac{h-1-\alpha}{h-1}$. Note that both inequalities hold for $\alpha > \frac{h-1}{2}$. We summarize the consequences for equilibrium behavior in the following lemma.

Lemma 4.6 *In a most responsive equilibrium of the open debate mechanism, players conceal information about the cost only if $\alpha < \frac{h-1}{2}$.*

The intuition for Lemma 4.6 is straightforward. If conflicts of interests are strong, high types are against implementation only if they know that costs are high, and low types are in favor of implementation only if they know that costs are low. Hence, if players do not learn information about the costs, they will not change their minds. Concealing information about the costs has no effect on individual votes and hence no effect on expected payoffs.

A similar line of reasoning applies to the non-disclosure of information about the benefits. Concealing b^l has an effect on a low type j's vote only if j reaches an information set $I_l(c^l)$, and j's belief at this information set satisfies $\mu_j(b^h) > \frac{\alpha}{(1-\alpha)(h-1)}$. The concealment affects a high type k's vote if k reaches an information set $I_h(c^h)$ and $\mu_k(b^h) > \frac{h-1-\alpha}{(1+\alpha)(h-1)}$ at this information set. Concealing b^h has an effect on a high type j's vote only if j reaches an information set $I_h(c^h)$, and j's belief in this information set satisfies $\mu_j(b^h) < \frac{h-1-\alpha}{(1+\alpha)(h-1)}$. The concealment affects a low type k's vote if k reaches an information set $I_l(c^l)$ and $\mu_k(b^h) < \frac{\alpha}{(1-\alpha)(h-1)}$ at this information set.

In a putative equilibrium in which no player conceals information about the benefits, no player has an incentive to deviate by concealing information about the benefits, if $\frac{1}{2} < \frac{\alpha}{(1-\alpha)(h-1)}$, and $\frac{1}{2} > \frac{h-1-\alpha}{(1+\alpha)(h-1)}$. Both inequalities are satisfied if $\alpha < \frac{h-1}{h+1}$.

Lemma 4.7 *In a most responsive equilibrium of the open debate mechanism, players conceal information about the benefits only if* $\alpha < \frac{h-1}{h+1}$.

For $\alpha \geq \frac{h-1}{2}$, there is no incentive to deviate from full information revelation. Hence, we are ready to identify the most revealing equilibrium for this parameter range, and to quantify induced the decision quality.

Lemma 4.8 *Consider the open debate mechanism. If* $\alpha \geq \frac{h-1}{2}$*, there is an equilibrium in which every committee member reveals the information he is endowed with. Low types vote against implementation unless they observe* b^h *and* c^l*. High types vote for implementation unless they observe* b^l *and* c^h.

Proof. There is obviously no profitable deviation in the voting stage. Revelation strategies follow from Lemmata 4.6 and 4.7. Q.E.D.

Proposition 4.2 *Consider the most revealing equilibrium of the open debate mechanism for* $\alpha > \frac{h-1}{2}$.
(i) In state (b^h, c^l)*, the probability of implementation is*

$$\frac{1}{2}\left(1+\delta^2\left(3\left(1-\frac{\delta}{2}\right)^2+\frac{1}{2}\left(\frac{\delta}{2}\right)^2\right)\right).$$

(ii) In state (b^l, c^h)*, the probability of rejection is*

$$\frac{1}{2}\left(1+\delta^2\left(3\left(1-\frac{\delta}{2}\right)^2+\frac{1}{2}\left(\frac{\delta}{2}\right)^2\right)\right).$$

(iii) Expected social surplus is

$$(h-1)\left(\delta^2\left(3\left(1-\frac{\delta}{2}\right)^2+\frac{1}{2}\left(\frac{\delta}{2}\right)^2\right)\right).$$

The proof of Proposition 4.2 can be found in Appendix B. Remember our initial assumption $\alpha < \frac{h-1}{h}$. Hence, the previously discussed parameter range is relevant only if $h \leq 2$. Note that in the equilibrium identified in Lemma 4.8, players are indifferent between sticking to their equilibrium revelation strategy and a possible deviation to conceal information which triggers a certain vote for (respectively against) implementation by the other types. Consider a low type's incentive to reveal c^l. As high types vote for implementation anyway (because they will not learn c^h), i can make sure that the project is implemented by voting for it if one of the other players reveals b^h. The probability of implementation will be the same, whether he conceals the information or he reveals it. Incentives to conceal information arise if the information makes the choice of the disliked alternative in the states in which players disagree more likely. If $\alpha < \frac{h-1}{2}$, there is an incentive for a low type to deviate from full information revelation. If the other players expect him to fully reveal his information, high types would vote against implementation in case only b^l is revealed in the communication stage. Hence, an equilibrium in which all players reveal their information exists only if $\alpha > \frac{h-1}{2}$.

We already observed that there may be an incentive for low types to conceal c^l if $\alpha < \frac{h-1}{2}$. The following lemma states that in fact both low types conceal c^l in the most revealing equilibrium for $\frac{h-1}{h+1} < \alpha < \frac{h-1}{2}$. A proof of the lemma can be found in Appendix B.

Lemma 4.9 *Consider the open debate mechanism and suppose* $\frac{h-1}{h+1} < \alpha < \frac{h-1}{2}$. *The most revealing equilibrium has the following properties. Low types conceal* c^l. *Each low type votes against implementation unless he observes* b^h *and* c^l. *High types vote for implementation unless they observe* b^l *and* c^h, *or if* b^l *is revealed in the communication stage, and (i) both low types have revealed* b^l, *(ii) only one low type revealed* b^l, *and* $\delta \leq \frac{h-1-2\alpha}{h-1-\frac{3}{2}\alpha}$, *or (iii) none of the low types revealed* b^l, *and* $\delta \leq \frac{\sqrt{h-1-\alpha}-\sqrt{\alpha}}{\sqrt{h-1-\alpha}-\frac{\sqrt{\alpha}}{2}}$.

Proposition 4.3 quantifies the highest decision quality which is attainable in equilibrium in the open debate mechanism for $\alpha \in \left[\frac{h-1}{h+1}, \frac{h-1}{2}\right]$.

Proposition 4.3 *Consider the most revealing equilibrium of the open debate mechanism for* $\frac{h-1}{h+1} < \alpha < \frac{h-1}{2}$.

(i) In state (b^h, c^l), the probability of implementation is

$$\tfrac{1}{2}\left(1+\delta^2\left(3\left(1-\tfrac{\delta}{2}\right)^2+\tfrac{1}{2}\left(\tfrac{\delta}{2}\right)^2\right)\right).$$

(ii) In state (b^l, c^h), the probability of rejection is

$$\begin{cases} \tfrac{1}{2}\left(1+\delta^2\left(\tfrac{13}{4}\left(1-\tfrac{\delta}{2}\right)^2+\tfrac{1}{2}\left(\tfrac{\delta}{2}\right)^2\right)\right), & \text{if } \delta > \delta'' \\ \tfrac{1}{2}\left(1+\delta\left(1+\tfrac{5}{4}\delta-2\delta^2+\tfrac{11}{16}\delta^3\right)\right), & \text{if } \delta' < \delta \leq \delta'', \text{ and} \\ 1-\tfrac{1}{2}\left(1-\left(\tfrac{\delta}{2}\right)\right)^4, & \text{if } \delta \leq \delta', \end{cases}$$

where $\delta' = \frac{\sqrt{h-1-\alpha}-\sqrt{\alpha}}{\sqrt{h-1-\alpha}-\frac{\sqrt{\alpha}}{2}}$, *and* $\delta'' = \frac{h-1-2\alpha}{h-1-\frac{3}{2}\alpha}$.

(iii) Expected social surplus is

$$\begin{cases} (h-1)\tfrac{1}{2}\delta^2\left(\left(\tfrac{\delta}{2}\right)^2+\tfrac{25}{4}\left(1-\tfrac{\delta}{2}\right)^2\right), & \text{if } \delta > \delta'' \\ (h-1)\tfrac{1}{2}\left(\delta^2\left(3\left(1-\tfrac{\delta}{2}\right)^2+\tfrac{1}{2}\left(\tfrac{\delta}{2}\right)^2\right)+\delta\left(1+\tfrac{5}{4}\delta-2\delta^2+\tfrac{11}{16}\delta^3\right)\right), & \text{if } \delta' < \delta \leq \delta'', \text{and} \\ (h-1)\tfrac{1}{2}\left(\left(\delta^2\left(3\left(1-\tfrac{\delta}{2}\right)^2+\tfrac{1}{2}\left(\tfrac{\delta}{2}\right)^2\right)\right)-\left(1-\left(\tfrac{\delta}{2}\right)\right)^4\right), & \text{if } \delta \leq \delta'. \end{cases}$$

For the purpose of this chapter it is not necessary to derive the set of equilibria for the entire parameter range. We now turn to the comparison of the two mechanisms, where we can already show for the previously discussed parameter range that the group debate mechanism may yield higher expected social surplus.

4.5 Comparison

If conflicts of interests are strong, i.e. $\alpha > \frac{h-1}{2}$, full information revelation is possible in the open debate (see Lemma 4.8). Hence, voters can base their decision on more information than in the group debate, and the states in which they agree will be identified with a higher probability. Therefore, the open debate mechanism performs strictly better than the group debate mechanism.

For smaller conflicts of interests, $\frac{h-1}{h+1} < \alpha < \frac{h-1}{2}$, in the group debate mechanism, high types vote against implementation if they observe low benefits and are uninformed about the costs. With open debate, low types have an incentive to conceal c^l. As a consequence, in case of a low type's silence, high types suspect him to conceal information. Then, they may refrain from voting against implementation although they know that benefits are low. This case arises if the probability with which a voter is endowed with information is not too low. If at the same time the probability of receiving information is low enough, it is likely that low types do not receive information, which will trigger the high types' suspicion in the open debate. Then, the group debate mechanism performs better in terms of the probability to reject the project in state (b^l, c^h). Overall, the group debate mechanism may yield a higher expected social surplus than the open debate mechanism. This is stated in the following proposition. A proof can be found in Appendix B.

Proposition 4.4 *Consider the case* $\frac{h-1}{h+1} < \alpha < \frac{h-1}{2}$ *and suppose* $\delta > \frac{h-1-2\alpha}{h-1-\frac{3}{2}\alpha}$*, and* $h < \frac{5}{3}$.
(i) There exists a δ^* *such that for* $\delta \in [\frac{h-1-2\alpha}{h-1-\frac{3}{2}\alpha}, \delta^*]$ *the probability of rejection in state* (b^l, c^h) *is higher in the group debate mechanism than in the open debate mechanism.*
(ii) There exist α^* *and* δ^{**}*, such that for all* $\alpha \in [\alpha^*, \frac{h-1}{2}]$*, expected social surplus is higher in the group debate mechanism than in the open debate mechanism if* $\delta \in [\frac{h-1-2\alpha}{h-1-\frac{3}{2}\alpha}, \delta^{**}]$.

4.6 Conclusion

In this chapter, we showed that it may be beneficial to restrict communication among decision makers with conflicting interests. When communication is allowed, committee members become suspicious when a member with different interests does not talk. The consequence is that agents react less to observable information than they would do if communication with committee members with conflicting interests was impossible. If decision makers possess information only with a small probability, the incidence that one of them does not talk arises with a high probability. Benefits from sharing (some) information are then outweighed by the inefficient usage of available information due to committee members' suspicion. If decision makers are restricted to communicate within groups of agents with aligned interests, a higher expected social surplus can be achieved.

We derived this result in a parsimonious model of a collective decision problem, which may be extended into several directions. In our model, it is crucial that communication between the groups is impossible. If decision makers could freely decide whom to inform, a decision maker would again make inferences when he is not informed by another agent. It would be interesting to see whether the group debate may endogenously arise in a model in which communication with other players is costly. Suppose for instance that a decision maker has to establish a costly link to another decision maker in order to send a message to him, and that this link has to be established before the decision makers are endowed with information. Another interesting extension would be to allow for endogenous information endowment. If information acquisition is costly, a decision maker's incentives to acquire information depend on with whom he can share the information, and whether other players will be listening.

Chapter 5

Speed and Quality of Collective Decision Making: Imperfect Information Processing

5.1 Introduction

In a seminal paper, Radner (1993) studies the efficient design of hierarchical structures when information processing takes time. Radner departs from the conventional assumption that individuals process information at infinite speed. He studies the problem of aggregating n data items at maximum speed when P information processors are available. This leads him to propose a hierarchical structure within which information is processed at maximum speed, the "reduced tree".[1] The virtue of the reduced tree is that processors on all levels simultaneously process information. This minimizes the delay of the entire information processing procedure. Radner's model can be applied to any information processing problem which requires repetitions of associative and commutative operations. One is the "max"-operation used in the collective decision problem which we consider in this chapter. Information processing in this case implies the pairwise comparison of possible alternatives and the identification of the best one.

Radner's analysis is focused on the efficient organization of information pro-

[1] The reduced tree will be described in greater detail in Section 5.2.3.

cessing with respect to three dimensions: (i) the size of the information processing task, (ii) the size of the organization, and (iii) the delay within which the task is completed. An organizational form is considered to be efficient, if given the number of processors involved, the delay cannot be reduced, and at the same time this delay cannot be achieved with a smaller number of processors.

In this paper, we add a new dimension to this evaluation of hierarchies: the quality of a decision. In Radner's original paper, which draws on a model brought forward by computer scientists (e.g. Gibbons and Rytter (1988)), it is assumed that processors work perfectly when they perform their task. But, in many real life situations individuals may make mistakes. In our analysis, we study a hierarchy which is composed of agents with imperfect calculation ability. Consequently, the evaluation of a hierarchy is carried out in terms of three dimensions: (i) the speed as well as (ii) the cost of information processing (i.e. the number of agents involved), and (iii) the quality of the decision.

We consider the project selection example proposed by Radner (1993). The organization's task is to select one item out of a class of n items. This corresponds, for example, to the choice of an investment project out of a set of competing investment opportunities.

Our mathematical analysis focuses on two measures: the probability that the best and the probability that the worst object is chosen by the hierarchy. We take these to be our measures of quality. Our main result is that reduced trees maximize quality for any number of data items, any number of processors and any potential for making mistakes. Thus, the reduced tree is efficient in terms of all three dimensions, speed, cost, and decision quality. Consequently, no trade-off exists between decision speed and decision quality in hierarchy design.

This chapter is related to a recent literature that extends Radner's framework into various directions. The reduced tree is designed for one-shot problems (to which we restrict attention). These are problems in which there is only one set of data to be processed, or in which the processing of the data is finished before another calculation task occurs. There are several contributions which assess the design of a hierarchy, or more generally a network of agents, when this is not the case (e.g. Van Zandt (1997, 1998); Meagher, Orbay and Van Zandt (2001)). Meagher and Van Zandt (1998) modify Radner's work with respect to the payment of managers. Orbay (2002) adds the frequency with which new data arrives as a new dimension to the analysis of efficient hierarchies. Prat (1997) studies hierarchies in which some managers are able to work faster than others and the wage a manager is paid is a function of his ability. It turns out that with these modifications – except for the one made by Prat (1997) – the reduced tree is still (close to) efficient.

To our knowledge, all previous models of time consuming information processing in hierarchies treat the information processing agents more or less like machines, which perfectly do what they are programmed to do. The value added of our work is to take into account the fact that human beings may make mistakes.[2]

In this regard, the problem studied in this chapter has certain similarities to that in Sah and Stiglitz (1986). Sah and Stiglitz assess the relative performance of two economic systems, namely a hierarchy and a polyarchy, when agents make mistakes in the assessment of projects. They find that a hierarchy is less likely to accept both bad projects and good projects than the polyarchy because the polyarchy gives a second chance to rejected projects and the hierarchy performs a second test on accepted projects. Sah and Stiglitz assume that agents use some benchmark for the assessment of projects and that they may implement as many projects as they like. Whereas in our model, the objective is to choose the best object out of a given set.

Another analysis of the quality of hierarchical decision processes has previously been carried out in Jehiel (1999). Jehiel considers the case where some information is lost during the aggregation procedure. With a certain probability, which depends on the size of a hierarchical unit, information aggregation within this unit is imperfect. Jehiel's measure of quality is the probability of perfect information aggregation. Optimal organizations consist of units with the same number of members.

Our formal analysis uses some results on the optimal (quality-maximizing) allocation of calculation tasks in van Zandt (2003). Van Zandt also studies the way in which the allocation of calculation tasks affects decision quality. In his setup, information aggregation is perfect, while the underlying information is changing over time. Delay generates costs because the decision becomes less and less appropriate over time. As in our model, this creates a cost when information

[2]One can argue that machines may make mistakes as well. In a companion paper, we study organization design in a moral hazard setup, where agents are free to decide whether to obey the program. Interestingly, the efficiency result extends to this setup.

is processed unequally. Van Zandt shows that the highest decision quality can be achieved with a balanced calculation tree, i.e. a sequence of calculation steps that guarantees a symmetric treatment of all objects.

The corresponding lemma and some further results on the optimal organization of calculation tasks will be briefly presented in Section 5.3 of this paper. In Section 5.4, we use these results to study the quality-maximizing structure of the hierarchical network, taking resource constraints into account. Based on the results on optimal calculation trees, we derive a set of necessary and sufficient conditions for maximum quality networks with a given number of processors and objects. We then show that these conditions are satisfied by all reduced trees. Since they also minimize delay, we conclude that reduced trees are efficient. Moreover, any outcome of an efficient organization can be achieved by a reduced tree.

5.2 The Model

The aim of this paper is to identify the quality-maximizing organization of information processing tasks for a given number of managers with imperfect information processing ability and to relate this design to the organizational form that completes the task fastest (Radner, 1993).

We first describe the decision problem and the limitations of agents' ability to deal with information. Next, we introduce our notion of hierarchies as well as their representation in trees and briefly recall Radner's (1992, 1993) result regarding the optimal organizational form with respect to the dimension of time.

5.2.1 The decision problem

We consider a decision problem in which one out of a set of n alternatives has to be chosen by a group of P identical agents. The alternatives (objects) are indexed by $i = 1, \ldots n$ and the agents (managers) are indexed by $p = 1, \ldots P$. The alternatives differ only with respect to quality. Ex ante (prior to information processing) the n objects are not distinguishable. However, there exists an objective ranking of the objects according to quality. We assume that quality is distributed in a such way that ties never occur, i.e. for two objects i and j, either i is of higher quality than j or vice versa. All managers have identical monotonous preferences regarding quality. In order to find the best alternative, the objects have to be compared. This is the information processing task that we study in this chapter. Agents are endowed with an inbox, a processing unit and a memory. In order to learn which alternative is the best one, an agent compares the objects pairwise. He reads an object from his inbox into his processing unit and compares it to the object in the memory. If the memory is empty, he stores the first object he reads without processing.

An information processing task, i.e. the comparison of two objects i and j, will be denoted $i \otimes j$. The result of the calculation, denoted $(i \otimes j)$, is meant to be the better of the two objects. However, information processing is imperfect: with probability $(1-q) < 1/2$, the agent makes a mistake[3] such that $(i \otimes j) = i$, although j is better than i. With probability q, no mistake is made and the agent correctly calculates $(i \otimes j) = j$.

Memory capacity is limited. In particular, an agent can store only one object. After having performed an operation, the agent stores the object he assesses to be the better one in his memory and takes the other one out of the set of possible alternatives.

Time enters the analysis in the following way: It takes a manager one unit of time to read an object he is supposed to process. We assume that neither the operation itself, nor sending a report (i.e. submitting a partial result to another agent) takes time. Thus, in one unit of time, a manager can perform the following tasks: (i) taking an object out of the inbox into the processing unit, (ii) comparing an object in the processing unit to the one in the memory (given that there are objects in both), and (iii) sending a message to the superior.

We do not introduce any assumptions about agents' preferences except for monotonicity in the chosen object's quality. Instead, we focus on the two extreme outcomes, which are of relevance for all quality distributions and for all von Neumann-Morgenstern utility functions: the event that the best object is chosen and the event that the worst object is chosen. Since it is the hierarchy's task to find the best alternative, it is natural to measure quality in terms of the probability that the best (worst) object will be the final outcome, i.e. in terms of the probability of success and the probability of complete failure. Accordingly, these probability measures quantify the quality of a decision in this chapter.[4]

[3]Another interpretation of the assumption that agents make mistakes is that they receive an imperfect signal about which object is better suited to fit their needs.

[4]In particular, these measures deliver a complete description of situations in which the

5.2.2 Trees

We restrict attention to hierarchical organizations. We follow Radner (1993) in defining a hierarchy (an organizational tree) as follows:

Definition 5.1 *A hierarchy $\mathcal{H}$ is a collection of objects (n data items and P managers), together with a relation among them, called "superior to". The relation has the following properties:*
1. Transitivity: If p is superior to p', and p' is superior to p'', then p is superior to p''.
2. Antisymmetry: If p is superior to p', then p' is not superior to p. p' is called p's subordinate.
3. There is exactly one object, called the root, that is superior to all the other objects.
4. Except for the root, every object has exactly one immediate superior.

Figure 5.1 illustrates an example for an organizational hierarchy $\mathcal{H}$ that processes 6 objects and the calculation tree $T(\mathcal{H})$ induced by the program.[5] In this figure (as in those following), $\boxed{p}$ represents manager p, and the objects are represented by their indices, $1, \ldots 6$. The "superior to" relation is represented by a link between the objects. A link from an object i to a manager p means that p reads object i, a link from a manager p to p' means that p reports the result of his calculation activities to p'.

The organizational hierarchy determines who performs which task and who reports to whom. We refer to the assignment of processing tasks as the program and call the organization a programmed hierarchy. The program gives rise to a calculation tree, illustrating the operations to be performed on the objects. There are n objects (represented as the leaves in the tree) to be processed, i.e. $(n-1)$ operations to be performed until the final result is obtained.

The program which underlies the trees depicted in Figure 5.1 gives the following instructions. In the first unit of time: agent 1, read object 1; agent 2, read object 3; agent 3, read object 5. In the second unit of time: agent 1, read 2, perform $1 \otimes 2$, report $(1 \otimes 2)$ to agent 4, and similar instructions for agents

hierarchy's task is to identify a certain object and there exists only one of its kind, e.g. a murderer or a thief (whom one would like to choose as a police department – in this case he represents the best object – and avoid choosing as a recruitment team).

[5] What we call a "calculation tree" is called a "binary tree" in Van Zandt (1993).

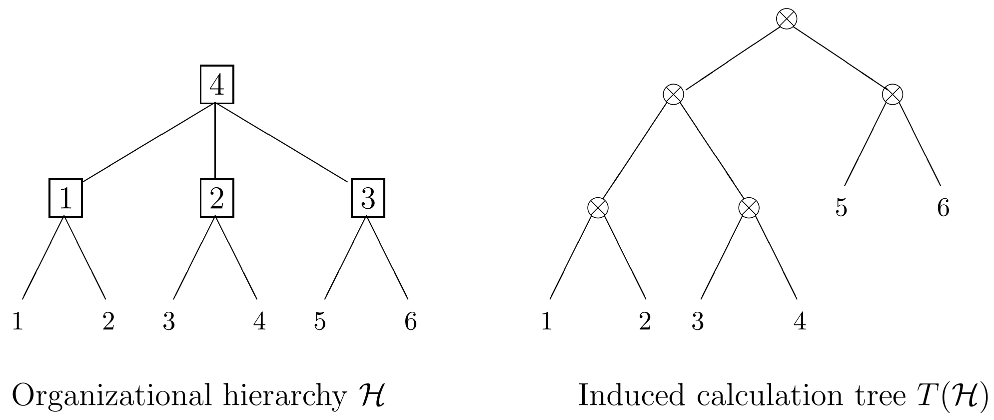

Figure 5.1: Organizational structure and induced calculation structure.

2 and 3. In the third unit of time: agent 4, read agent 1's report, in the fourth unit of time, read 2's report, perform $(1 \otimes 2) \otimes (3 \otimes 4)$, and in the fifth unit of time, read agent 3's report, perform $((1 \otimes 2) \otimes (3 \otimes 4)) \otimes (5 \otimes 6)$, and report the result.

Note that a calculation tree can be induced by different organizational structures. For example, the calculation tree in Figure 5.1 is also obtained if agent 4's tasks are performed by one of the managers on the lower hierarchy level. However, the organizational design matters in terms of speed.

In our setting with imperfect information processing ability, the associative law of binary operations does not hold anymore. This is why quality depends on the order of calculations. What matters in terms of quality is the calculation tree. The particular organizational tree by which it is induced does not matter in our setup as managers are assumed to be identical with respect to their calculation ability. Hence, who performs a particular operation is irrelevant for quality. However, feasibility restrictions with respect to the hierarchical organization (as the number of managers) limit the set of calculation trees that can be implemented.

5.2.3 Radner's efficiency result

In this section, we briefly recall Radner's (1993) results concerning the optimal organizational design. Decision delay is affected by the extent of parallel computation, i.e. how much information is handled simultaneously. Radner (1992, 1993) derives the following delay-minimizing organization of information processing: (i)

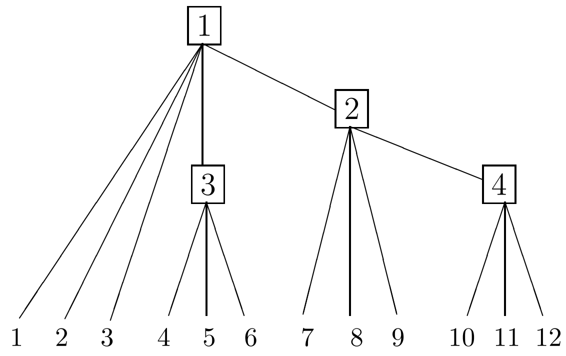

Figure 5.2: Reduced tree.

Number the P managers subsequently from 1 to P and assign n/P objects to each manager. (If n/P is not an integer, assign $\lfloor n/P \rfloor$ to each manager and another one to the last $n \bmod P$ ones.)[6] (ii) Assign manager i's partial result to $(i-1)$, for each even i. Manager $(i-1)$ therewith becomes i's immediate superior and i is $(i-1)$'s subordinate (if P is odd, one manager remains unconnected). (iii) Rename the managers who are not yet somebody's subordinate, assigning the number 1 to the manager with the largest number. Repeat (ii) and (iii) until a single manager remains (the top manager).[7] Figure 5.2 depicts a reduced tree in which 4 managers process 12 objects.

The virtue of this design is that information is processed in parallel as far as it is possible. The n objects are aggregated to the final result in $\lfloor n/P \rfloor + \lceil \log_2 (P + n \bmod P) \rceil$ units of time.[8] Note that in order to maximize the speed of information processing, the efficient number of managers never exceeds $\lfloor n/2 \rfloor$.

5.3 Calculation trees with maximum quality

In this section, we study the optimal organization of calculation tasks without paying attention to feasibility constraints with respect to the hierarchical form by which it could be induced. We are interested in the decision quality associated

[6] $\lfloor x \rfloor$ denotes the largest integer $\leq x$, and $\lceil x \rceil$ is the smallest integer $\geq x$.

[7] This construction slightly deviates from the one proposed by Radner (1993), which does not have an effect on speed, but – as will become clear later on – may affect quality.

[8] See Radner (1993) for the proof that this is indeed the delay minimizing organization.

with the organization of information processing.[9]

We take a calculation tree T as given. The calculation tree determines a calculation path for each object, and in particular how many comparisons an object has to pass before being selected. Let $\delta(i,T)$ be the length of the path from leaf i to the root in tree T,[10] and let $\delta^T = (\delta(1,T), ..., \delta(n,T))$. Given that ex ante the objects are of the same expected quality, the decision quality associated with T is fully described by δ^T.

In order to choose the best object, all calculations performed on this object have to be correct. To choose the worst object, all operations performed on it have to entail mistakes. Since the combinatorics for choosing the best or the worst item out of the given set are the same, the calculations are analogous. To simplify the exposition, we focus on the probability of choosing the best object. All of our results apply to quality measured in terms of the probability of picking the worst object as well. This can easily be seen by replacing q with $(1-q)$ in the calculations and is therefore omitted.[11]

If leaf i is the best object and if i is supposed to be processed $\delta(i,T)$ times in tree T, it will be chosen with probability $q^{\delta(i,T)}$. Object i is the best object with probability $1/n$. Hence, our quality measures can be defined as follows.

Definition 5.2 $\pi(x,n,T) = \frac{1}{n}\sum_{i=1}^{n} x^{\delta(i,T)}$.

Lemma 5.1 *With calculation tree T, the probability of choosing the best object is $\pi(q,n,T)$, and the probability of choosing the worst object is $\pi(1-q,n,T)$.*

Let $\mathcal{T}$ denote the set of calculation trees with n leaves. With respect to decision quality, we use the following optimality criterion.

Definition 5.3 *A calculation tree is optimal for the processing of n objects, if it solves* $\max_{T\in\mathcal{T}} \pi(q,n,T)$ *and* $\min_{T\in\mathcal{T}} \pi(1-q,n,T)$.

[9] A different problem with a similar mathematical representation is studied in van Zandt (2003). We thank a referee for pointing out to us that van Zandt's results might help us to generalize our earlier (Schulte and Grüner, 2004) results on the comparison of different hierarchies.

[10] When there is no ambiguity, we will suppress the argument T and use the notation δ_i.

[11] Note that it is not universally valid that an organization which chooses the best object with a higher probability than another organization also has a lower chance to pick the worst object (see Schulte and Grüner (2004)). However, in our model, this is a feature of the optimal organization.

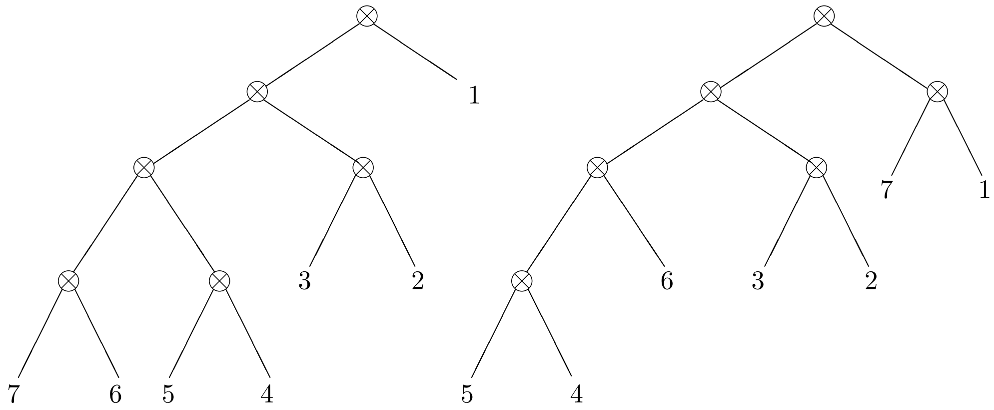

Figure 5.3: Calculation tree before and after Manipulation $(*)$.

Van Zandt (2003) has shown in a different context that calculation trees in which the distance from a leaf to the root differs too much among the leaves are dominated by trees in which these distances are more equal. The same is true in our setup, as we will show in the following. Consider an arbitrary calculation tree T, characterized by δ^T. Arrange the tree such that it starts with the deepest nodes (as shown in Figure 5.3), and label the leaves in descending order as they appear in the tree. We make use of a quality enhancing reorganization of the leaves in order to derive the optimal calculation tree.

Manipulation $(*)$, "switching" (van Zandt, 2003) :

1. Delete operation $n \otimes (n-1)$ and perform the successive operations in the affected branch on $(n-1)$ instead of $(n \otimes (n-1))$.
2. Add operation $1 \otimes n$.

Call the new tree T'. Manipulation $(*)$ is illustrated in Figure 5.3.

This manipulation has the following effects:

1. object 1 was formerly processed $\delta(1,T)$ times, now it is processed $\delta(1,T)+1$ times,

2. object n was formerly processed $\delta(n,T)$ times, now it is processed $\delta(1,T)+1$ times,

3. object $(n-1)$ was formerly processed $\delta(n,T)$ times, now it is processed $\delta(n,T)-1$ times.

Thus, the effect of the manipulation on the probability of choosing the best object is:

$$\begin{aligned}\pi\left(q,n,T\right)-\pi\left(q,n,T'\right) &= \tfrac{1}{n}\left(-2q^{\delta(n,T)}-q^{\delta(1,T)}+2q^{\delta(1,T)+1}+q^{\delta(n,T)-1}\right)\\ &= \tfrac{1}{n}\left(2q-1\right)\left(q^{\delta(1,T)}-q^{\delta(n,T)-1}\right). \qquad (5.1)\end{aligned}$$

This expression is positive, if $\delta(1,T)<\delta(n,T)-1$. Hence, the proposed manipulation increases the probability of choosing the best object as long as this inequality holds. If we replace q with $1-q$ in (5.1), the expression is negative, verifying that Manipulation $(*)$ leads to a smaller probability of choosing the worst outcome as well.

The intuition for this effect is the following. In the original tree, object 1 had the best position concerning the probability of being chosen, because it had to survive the smallest number of comparisons. Object 1 is the best (worst) object with probability $1/n$. Manipulation $(*)$ replaces object 1 with an object which has a higher chance to be the best one and a lower chance to be the worst one. This is so because processing a subset of items always increases quality as compared to randomly choosing an object. In the other affected branch, the partial result $(n\otimes n-1)$ is replaced with object $(n-1)$, to which a lower probability is attached that it is the best object and a higher probability that it is the worst one. But this object's position is the worst one in the calculation tree concerning the probability of being chosen. Hence, the first effect dominates the latter.

We now show that by repeatedly applying Manipulation $(*)$ (with appropriate renaming of the objects) the best (highest quality) calculation tree is obtained. Each manipulation increases (decreases) the probability of choosing the best (worst) object, as long as $\delta_1<\delta_n-1$. One can stop this procedure as soon as $\delta_1\geq\delta_n-1$, i.e. when the objects are processed as equally as possible. Equation (5.1) implies that the calculation tree cannot be improved further with this manipulation. In particular, applying Manipulation $(*)$ again only changes the position of the objects in the tree, but does not affect the probability of choosing the best object. Following van Zandt (2003), we call the resulting tree a *balanced tree* (see Figure 5.4).

Definition 5.4 *Consider a calculation tree T with n leaves. Name the leaves*

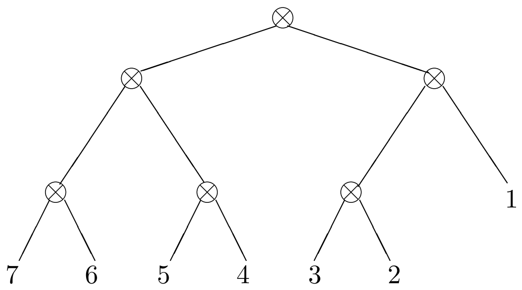

Figure 5.4: Balanced tree.

$i = 1, \dots n$ such that $\delta_1 \leq \delta_2 \leq ... \leq \delta_n$, where δ_i is the length of the path from i to the root. A tree with the property $\delta_1 \geq \delta_n - 1$ is called a balanced tree.

Is there now an alternative manipulation that enhances decision quality even further? Assume that this is the case. Then either (i) the manipulation again leads to a balanced tree, or (ii) it yields a tree that is not balanced. In case (i), the manipulation does not increase decision quality because all balanced trees (given n and q) produce the same quality. In case (ii), the algorithm described above can be applied to the resulting non-balanced tree, resulting in a balanced tree which is superior in terms of quality. This implies that the resulting balanced tree is better than the original balanced tree, and hence a contradiction.

This yields Lemma 5.2.

Lemma 5.2 (van Zandt, 2003) *Consider a calculation task on n objects. All optimal calculation trees are balanced trees, i.e. $\delta_i - \delta_j \leq 1$, for all objects i and j.*

We can also reverse the argument applied in this section. It is straightforward that if one applies a manipulation converse to Manipulation $(*)$, quality decreases. By repeatedly doing so, one gets a calculation tree with the property $\delta_1 = 1,\ \delta_i = \delta_{i-1} + 1$ for all $i \neq 1, n$, and $\delta_n = \delta_{n-1}$. Call this tree a *serial tree.*

Corollary 5.1 *Consider a calculation task on n objects. Serial processing yields the lowest quality.*

The serial tree and serial subtrees will play an important role in the our analysis of efficient organizations. Two further useful results are stated in Lemmata 5.3 and 5.4.

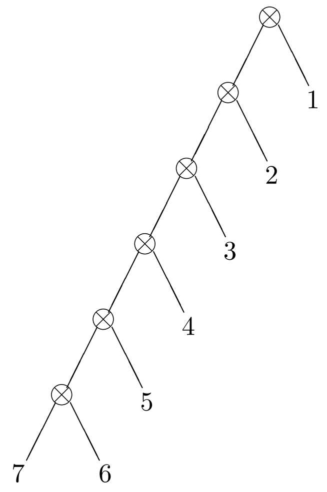

Figure 5.5: Serial tree.

Lemma 5.3 (van Zandt, 2003) *In a serial tree S with n leaves, the probability of choosing the best object is*

$$\pi(q, n, S) = \frac{1}{n}\frac{q}{1-q}\left(1 + q^{n-2}(1-2q)\right).$$

To save on notation we introduce Definition 5.

Definition 5.5 $s(x) = \frac{q}{1-q}\left(1 + q^{x-2}(1-2q)\right)$.

Consider a partial result, produced by serially processing a subset of n_p (out of n) items, e.g. the result of the first three operations in Figure 5.5, $(((7 \otimes 6) \otimes 5) \otimes 4)$.

Lemma 5.4 *(i) The probability that the partial result produced in a serial subtree containing n_p raw data items is the best object is $(1/n)s(n_p)$.*
Moreover, (ii) for $x > 0$

$$s(n_p + x) = s(n_p) + q^{n_p - 1}\left(s(x+1) - 1\right).$$

Proof. Part (i) follows from Lemma 5.3. To see (ii), note that in two serial subtrees, one containing n_p items, and the other one containing $n_p + x$ items, where $x > 0$, the $n_p - 1$ items processed last have the same probability of reaching the root in both trees. The item processed first in the smaller tree is the best object with probability $1/n$. This object is replaced in the larger tree with a partial result (produced within another serial subtree) which is the best object with probability $(1/n)s(x+1)$. Q.E.D.

5.4 Efficient organizations

We use the following efficiency criterion in our analysis.

Definition 5.6 *A hierarchy is efficient if no alternative organization exists for processing a given set of n objects which performs better on one of the dimensions (i) quality, (ii) speed, and (iii) cost, and at least equally well on the other dimensions.*

In this section, we show that for a given information processing problem and a given number of managers P, reduced trees maximize quality. Since the reduced tree is also the fastest way to deal with n objects, this organization is efficient.

The efficiency of the reduced tree can directly be established for an organization with $\lfloor n/2 \rfloor$ managers. Recall the properties of a reduced tree from Section 5.2.3. In the first phase of information aggregation, each manager aggregates the raw data items assigned directly to his inbox. A manager's processing activity in this phase can be depicted by a serial tree. With $\lfloor n/2 \rfloor$ managers, each of the serial trees has two leaves.[12] Moreover, the reporting structure induces a balanced calculation tree, where the "leaves" are the results of the first aggregation phase. Hence, a reduced tree with $\lfloor n/2 \rfloor$ managers induces a balanced calculation tree.

Proposition 5.1 *A reduced tree with $\lfloor n/2 \rfloor$ managers is efficient.*

Consequently, there is a non-trivial upper bound for the size of efficient hierarchies.

Corollary 5.2 *No efficient hierarchy employs more than $\lfloor n/2 \rfloor$ managers.*

We now turn to a complete characterization of efficient hierarchies. For this purpose, it is useful to distinguish between an object that has been compared to another one and an object that has not yet been processed. We call the former a *partial result* and the latter a *raw data item*. Let n_p denote the number of raw data items p is supposed to handle, with $\sum_{p=1}^{P} n_p = n$. If fewer than $\lfloor n/2 \rfloor$ managers are available, there must be more than one manager p such that $n_p > 2$. In a hierarchy with fewer than $\lfloor n/2 \rfloor$ managers, the balanced calculation tree is

[12] If n is odd, one of the trees has three leaves.

in general not implementable. The reason is that limited memory capacity forces each manager to process the information he is supposed to handle serially.

We now derive the quality-maximizing hierarchy design with fewer than $\lfloor n/2 \rfloor$ managers. As it turns out, reduced trees are optimal. To gain an intuition why this is the case, consider Figure 5.6, which represents the calculation tree induced by the reduced tree depicted in Figure 5.2.

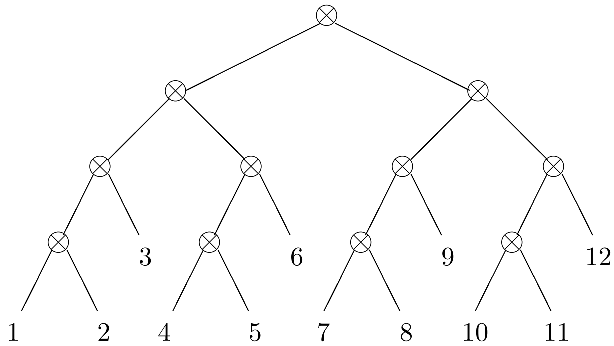

Figure 5.6: Calculation tree induced by the reduced tree in Figure 5.2.

The calculation tree contains four serial subtrees of equal length, hence any partial result is the best object with equal probability. We know from the previous section that there is no better way to organize the processing of the partial results than a balanced tree. Moreover, given that the partial results are processed in a balanced tree, i.e. all partial results have the same distance m_p to the root, a reassignment of raw data items cannot increase quality: Consider a reassignment of one object from one agent to another. This reorganization affects two serial subtrees: Both have the same number n_p of leaves initially. After the reassignment, one of them has $n_p - 1$ leaves, and the other one has $n_p + 1$ leaves. The effect on $\pi(.)$ is

$$\begin{aligned} \Delta\pi(.) &= \tfrac{1}{n} q^{m_p} \left(s(n_p - 1) + s(n_p + 1) - 2s(n_p) \right) \\ &= \tfrac{1}{n} q^{m_p} (1 - 2q)\, q^{n_p - 1} (1 - q) < 0 \Leftrightarrow q > \tfrac{1}{2}. \end{aligned}$$

However, the reasoning above does not tell us anything about the reduced tree's performance relative to hierarchies in which neither every agent processes the same number of raw data items, nor the partial results are aggregated in a balanced tree. Our strategy to exclude such structures is to consider an arbitrary hierarchy $\mathcal{H}$ and to look for a feasible reorganization of calculation tasks such that quality increases. If we find such a reorganization, hierarchy $\mathcal{H}$ is not optimal with respect to quality. This yields necessary conditions for optimal hierarchy design. Step by step, we restrict attention to hierarchies in which the necessary conditions are met. We finally show that the set of necessary conditions is sufficient.

Consider a programmed hierarchy $\mathcal{H}$ with $P < n/2$ managers and its calculation tree $T(\mathcal{H})$. Let $n_p \geq 0$ be the number of raw data items processed by agent p and let $r_p \geq 0$ be the number of reports processed by agent p. If p' reports to p, we denote by $r_{p'p}$ the report p' sends to p. We may restrict attention to the case that p' processes strictly more than one object, since otherwise $\mathcal{H}$ could be reprogrammed assigning the object directly to p and firing p'. We denote by $\delta(i, T(\mathcal{H}))$ the length of the path from leaf i to the root in the calculation tree induced by $\mathcal{H}$, and with $\delta(r_{p'p}, T(\mathcal{H}))$ the length of the path from report $r_{p'p}$ to the root. The first two necessary conditions on quality-maximizing hierarchies are provided in Lemmata 5.5 and 5.6.

Lemma 5.5 *In a quality-maximizing hierarchy with $P < n/2$ managers, every agent p, for whom $n_p > 0$ and $r_p > 0$, processes the raw data items before processing a report.*

Proof. Suppose p was processing report $r_{p'p}$ before being finished with the raw data. Then there exists an object i such that $\delta(i, T(\mathcal{H})) < \delta(r_{p'p}, T(\mathcal{H}))$. The probability that i is the best object is $1/n$, whereas $r_{p'p}$ – containing more than one raw data item – is the best object with probability higher than $1/n$ and the worst one with probability lower than $1/n$. Exchanging the positions of object i and report $r_{p'p}$ increases the probability that report $r_{p'p}$ will be chosen from $q^{\delta\left(r_{p'p}, T(\mathcal{H})\right)}$ to $q^{\delta(i, T(\mathcal{H}))}$ (if it is the best object), and vice versa for object i. Hence, quality increases. Q.E.D.

A consequence of the lemma above is that for any agent p with $n_p > 0$ the calculation subtree describing this agent's task contains a serial subtree connecting the n_p raw data items. The reorganization of the hierarchy proposed in the proof of Lemma 5.5 and the associated effect on the calculation tree are illustrated in Figure 5.7.

Lemma 5.6 *In a quality-maximizing hierarchy with $P < n/2$ managers, every agent processes raw data.*

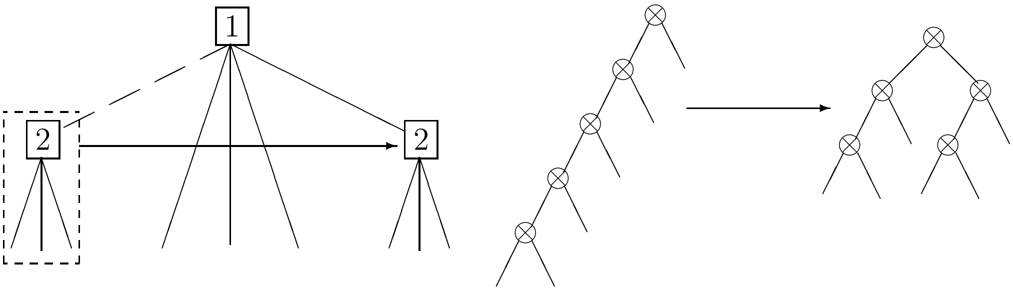

Figure 5.7: Process raw data first. *(Lemma 5.5)*

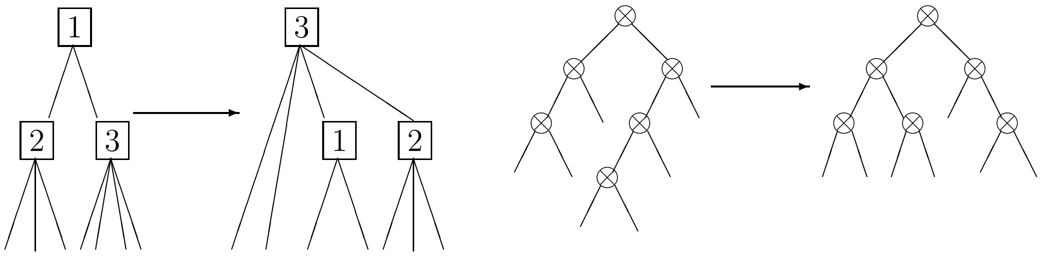

Figure 5.8: All agents process raw data. *(Lemma 5.6)*

Proof. Suppose $n_p = 0$ for some agent p. Assign p's tasks to one of his direct subordinates and preserve the order of calculations in this subtree. Note that quality is unaffected. Look for the agent p' who processes the largest amount of raw data. If $n_{p'} > 3$, assign two items to p and let p report the result to p'. Let p' read the report after finishing his raw data. Note that this yields a manipulation of the calculation tree in the spirit of Manipulation $(*)$ and hence increases quality. If $n_{p'} = 3$, search for another agent p'' processing three raw data items. Assign an item from each of these managers to p. Denote by $N_{p'p''}$ the set of remaining raw data items processed by p' and p''. Identify the agent who processes the object i, for which $\delta(i, \mathcal{H}) = \min \{\delta(j, \mathcal{H})\}_{j \in N_{p'p''}}$, and let p report to this agent. Let p's report be read immediately after the raw data. Again, the reorganization yields a manipulation of the calculation tree in the spirit of Manipulation $(*)$ and hence increases quality. Q.E.D.

The intuition for the lemma above is straightforward: By involving more processors in the processing of raw data items, serial subtrees can be transformed into "more balanced" ones. We know from the previous section that this enhances quality. The reorganization of the hierarchy is illustrated in Figure 5.8.

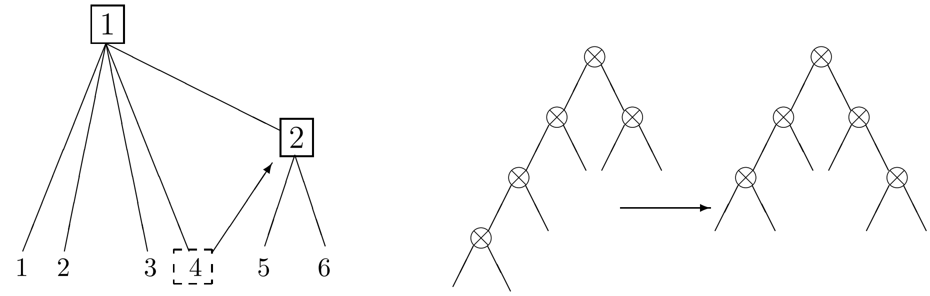

Figure 5.9: Reorganization used to prove Lemma 5.7.

From Lemmata 5.5 and 5.6 we know that every quality-maximizing hierarchy induces a calculation tree containing P serial subtrees. It remains to be studied (i) how many leaves (n_p) these subtrees optimally have, (i.e. how many raw data items shall be assigned to each manager), and (ii) how the serial subtrees are optimally arranged, (i.e. the optimal reporting structure). Note that with regard to the arrangement of the P serial subtrees, any tree can be implemented, because we are endowed with P processors, and there are only $P-1$ operations left to perform.

Let m_p denote the distance from p's result of the raw data processing phase to the root.

Lemma 5.7 *In a quality-maximizing hierarchy,* $(m_p + n_p) - (m_{p'} + n_{p'}) \leq 1$ *for all* p *and* p'.

Proof. The probability of choosing the best object is $\pi(.) = \sum_{p=1}^{P} \frac{1}{n} q^{m_p} s(n_p)$. Consider the effect of a reassignment of a raw data item from p to p'. We have:

$$\begin{aligned}
\Delta\pi(.) &= (1/n)(q^{m_p}(s(n_p-1) - s(n_p)) + q^{m_{p'}}(s(n_{p'}+1) - s(n_{p'}))) \\
&= (1/n)(q^{m_p}(1-2q)q^{n_p-2} + q^{m_{p'}}(2q-1)q^{n_{p'}-1}) \\
&= (1/n)(2q-1)(q^{m_{p'}+n_{p'}-1} - q^{m_p+n_p-2}) \\
&> 0 \quad \Leftrightarrow m_p + n_p > m_{p'} + n_{p'} + 1.
\end{aligned}$$

Hence, we can increase quality using the proposed reorganization as long as there exist p and p' such that $(m_p + n_p) - (m_{p'} + n_{p'}) > 1$. Q.E.D.

According to Lemma 5.7, the distance from a serial subtree's deepest leaf to the root of the calculation tree differs by at most one unit among all serial subtrees. The last step is to determine the optimal reporting structure. The optimal assignment of raw data items then follows from Lemma 5.7.

Lemma 5.8 *In a quality-maximizing hierarchy, for all* p *and* p'*, (i)* $m_p - m_{p'} \leq 1$*, and (ii)* $n_p - n_{p'} \leq 2$*, where* $n_p - n_{p'} = 2$ *only if* $m_{p'} - m_p = 1$*.*

Proof. (i) Name the agents such that $m_1 = m_2 \geq ... \geq m_P$. If $m_p = m_{p\prime}$, and $n_p < n_{p\prime}$, name p and p' such that p gets the lower number. Hence, agent 1 processes the lowest number of raw data items, and his partial result has the longest distance to the root. Agent P's partial result has the shortest distance to the root. Note that the serial subtree representing agent 1's raw data processing activities must be connected directly to another serial subtree. That is, there exists an agent p' such that either agent 1 or agent p' is engaged only in raw data processing, and the other agent reads the report immediately after his raw data items.

Now, assume that $m_p - m_{p'} \leq 1$ does not hold for all p, i.e. assume that there exists an integer $k > 1$ such that $m_1 = m_P + k$. We show that in this case the following reorganization enhances quality: Reassign a raw data item from agent P to agent 1. Let agent 1 report the result of his raw data processing activities to agent P, and let agent P read this report after finishing his remaining raw data processing activities. Assign agent 1's remaining tasks (if there are any) to agent p'. Note that the raw data items processed by p' move up one position in the calculation tree when removing agent 1's partial result.

The effect of this reorganization on quality is the following:

$$\Delta\pi(.) = (1/n)\{q^{m_1-1}((1-q)s(n_{p\prime}) - qs(n_1)) \tag{5.2}$$
$$+q^{m_P+1}(s(n_1+1)-1)\}. \tag{5.3}$$

The effect of removing the serial subtree representing agent 1's raw data processing activities is captured by (5.2), and the effect of reattaching it is summarized in (5.3). Note that the serial subtree representing agent 1's tasks replaces the raw data item processed last by agent P in the original program. Lemma 5.7 requires that $n_{p'} \in \{n_1,\ n_1+1\}$. Let D assume the value 1 if $n_{p'} = n_1 + 1$, and 0 otherwise. We have:

$$\begin{aligned}\Delta\pi(.) &= (1/n)\{q^{m_P+k-1}((1-q)s(n_1+D) - qs(n_1)) \\ &\quad +q^{m_P+1}(s(n_1+1)-1)\}.\end{aligned}$$

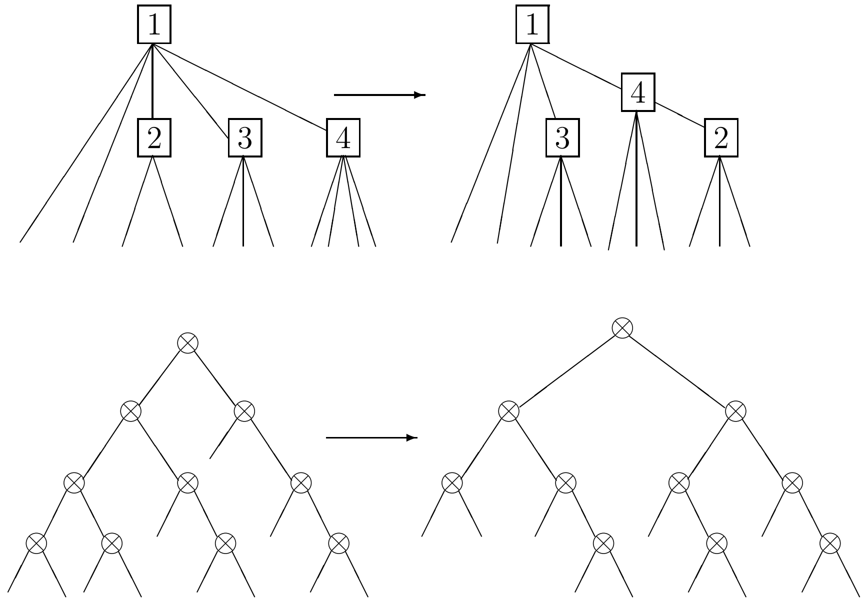

Figure 5.10: Reorganization used to prove Lemma 5.8.

$$\begin{aligned}
& \Delta\pi(.) > 0 \\
\Leftrightarrow\ & q^{k-2}((1-2q)s(n_1) + Dq^{n_1-1}(2q-1)(1-q)) + qs(n_1) + q - 1 > 0 \\
\Leftrightarrow\ & (1-q)(\tfrac{q}{1-q}(1+(1-2q)q^{k-3})s(n_1) - 1 + Dq^{n_1+k-3}(2q-1)) > 0 \\
\Leftrightarrow\ & (1-q)(s(k-1)s(n_1) - 1) + Dq^{n_1+k-3}(2q-1)(1-q) > 0
\end{aligned} \tag{5.4}$$

The second term in (5.4) is either zero, or positive if $q > 1/2$, and negative if $q < 1/2$. The first term is positive if and only if $s(k-1)s(n_1) - 1$ is positive. Note that $s(x) > 1$ for $x \geq 1$ if and only if $q > 1/2$. Hence, $\Delta\pi(.) > 0$ if and only if $q > 1/2$.

Part (ii) follows from Lemma 5.7 and part (i). Q.E.D.

Note that, if $n_p - n_{p'} = 2$ and $m_{p'} - m_p = 1$, a reassignment of one raw data item from p to p' yields a permutation of δ^T, but has no effect on quality. That is, a quality-maximizing hierarchy which has the property $n_p - n_{p'} = 2$, for some p and p', coexists with another quality-maximizing organization with the property $n_p - n_{p'} \leq 1$, for all p and p'.

Part (i) of Lemma 5.8 states that an optimal hierarchy has a reporting structure such that the P partial results of the first processing phase are aggregated in a balanced calculation tree. Part (ii) requires an equal assignment of the raw data items (up to "integer leftovers"). We are in a position to state our main results.

Proposition 5.2 *A hierarchy with P managers maximizes quality if and only if it has the properties stated in Lemmata 5.5-5.8.*

Proof. Necessity has been established already. To verify sufficiency, note that any feasible reorganization has either no effect on quality or yields a violation of at least one of the necessary conditions stated in Lemmata 5.5-5.8. Q.E.D.

Proposition 5.3 *All reduced trees are efficient.*

Proof. Reduced trees have the properties described in Lemmata 5.5-5.8, hence they maximize quality given P. We know from Radner (1993) that a reduced tree achieves the minimum delay for a given number of processors. Q.E.D.

As a reduced tree maximizes the speed of the decision procedure given the number of managers, as well as the quality of the decision, no slower working hierarchy with the same number of managers can be efficient.

Corollary 5.3 *A hierarchy with P managers processing n objects is efficient only if it achieves minimum delay. Any outcome of an efficient organization can be achieved by a reduced tree.*

5.5 Conclusion

We studied a problem of efficient decentralized information aggregation in a setup with information processing imperfections. Our results indicate that a hierarchy designer does not face a trade-off between the speed and the quality of information aggregation. The reduced tree proposed by Radner (1993) is a hierarchy in which information processing takes minimum time and delivers maximum quality for a given number of processors and objects. Moreover, we find that highest quality can be obtained by the reduced tree with $\lfloor n/2 \rfloor$ managers (which is again the fastest of its class).

There are several useful extensions of our framework. First, it would be desirable to consider more general measures of decision quality. This would require the specification of a von Neumann-Morgenstern utility function, as well as assumptions on the distribution of quality. Our results hold regardless of the form the utility function takes or how quality is distributed.

As a referee pointed out to us, processing imperfections may have an impact on the properties of the efficient set of hierarchies only if the associative law of binary operations does not hold anymore. For combinations of calculation tasks and processing imperfections other than the ones considered in our paper, it may still hold. It is worth studying which combinations yield a violation of the associative law and whether alternative specifications would affect our efficiency results.

We introduced an information processing imperfection into the analysis of decentralized information aggregation by restricting the calculation ability of agents in an intuitive, but rather simplifying manner. In our setup, agents' mistakes do not depend on the quality of the two compared items. Intuitively, mistake making should depend on the task to be performed. To incorporate this consideration into our model, one could make use of probabilistic choice models, such as Luce (1959). Again, this modification would require the specification of the quality distribution and of the utility function.

One may also allow agents to influence the individual probability of making a mistake through effort. This issue is addressed in Grüner and Schulte (2004), where a game-theoretical approach is taken to study the incentives for effort provision in carrying out a decentralized information processing task.

Finally, we restricted attention to hierarchical organizations. In such an organization, calculations cannot be repeated and agents make one final report to their superior. If information processing is imperfect, one would like to repeat calculations to increase quality. However, the hierarchical structure precludes the possibility of sending an object to multiple processors. Whether or not a more general structure is feasible depends on the agents' information-storing abilities and on the information structure. For instance, if the objects cannot be copied, and only one agent can work on an object at any time (e.g. job candidates who need to be interviewed), then it is impossible to send objects to multiple processors and hierarchies are the only feasible organization in a setting with limited memory capacity. In our notion of a hierarchy, we allowed for only one upward link for each manager. However, one might want to allow a manager to send more than one report to his superior. With respect to the speed of information processing, it is optimal to implement a hierarchical structure as defined in this paper. In terms of quality, it may make sense to allow for multiple reports to get closer to the balanced calculation tree. Concerning these modifications, one would have to specify a different information processing technology. Both modifications – multiple reports and repeated calculations – would involve a trade-off between speed and quality that does not play a role in our framework.

Appendix A

Appendix to Chapter 3

Derivation of Equation (3.2)

Given k_α α-signals and k_β β-signals, Bayesian updating yields:

$$\begin{aligned} p(\omega = A|\kappa = k_\alpha - k_\beta) &= \frac{\frac{1}{2}q^{k_\alpha}(1-q)^{k_\beta}}{\frac{1}{2}q^{k_\alpha}(1-q)^{k_\beta} + \frac{1}{2}q^{k_\beta}(1-q)^{k_\alpha}} \\ &= \frac{q^{k_\alpha - k_\beta}}{q^{k_\alpha - k_\beta} + (1-q)^{k_\alpha - k_\beta}} \end{aligned}$$

Defining $k = k_\alpha - k_\beta$ gives us:

$$p\left(A\left|\kappa = k\right.\right) = \frac{q^k}{q^k + (1-q)^k}.$$

Derivation of Equation (3.3)

There are exactly k more α-signals than β-signals (k possibly negative) within the group of voters except for i if there are exactly $\frac{n-1+k}{2}$ α-signals, and (the residuum) $n - 1 - \frac{n-1+k}{2}$ β-signals. We have to sum up all these cases:

$$\mu_i(\kappa_{-i} = k) = \begin{cases} \sum\limits_{J \subseteq \{1,\ldots n\}\setminus\{i\}:|J| = \frac{n-1+k}{2}} \prod\limits_{j\in J} \mu_i(\sigma_j = \alpha) \prod\limits_{l \in \{1,\ldots n\}\setminus\{J\cup i\}} \mu_i(\sigma_l = \beta), \\ \qquad \text{for } k \in E\{-n+1, \ldots, n-1\} \\ 0, \qquad \text{for } k \notin \{-n+1, \ldots, n-1\}. \end{cases}$$

Proof of Proposition 3.3

First note that in a unanimous voting strategy profile (conditional on the communication outcome) as in the potential equilibrium, no single vote has an effect on the outcome. Hence there exists no profitable deviation at the voting stage. Moreover, the voting strategy profile has the property that no private information (information which was not revealed in the communication stage) will be aggregated in the voting stage. The collective decision depends only on the evidence presented in the communication stage: $x = a$, if $k(c) \geq k + 1$, and $x = b$ else.

In the following, we can take the decision rule as given. Note that the revelation of an α-signal can only have the effect to change the decision from b to a, and vice versa for the revelation of a β-signal.

Agent i has an incentive to reveal an α-signal if and only if he believes that $p_i(\omega = A) \geq \theta_i$ conditional on the event that his revelation changes the decision from b to a, i.e. conditional on the evidence being k without his revelation. He has an incentive to reveal a β-signal if and only if he believes that $p_i(\omega = A) \leq \theta_i$ conditional on the event that his revelation changes the decision from a to b, i.e. conditional on the evidence being $k + 1$ without his revelation.

γ^* is a communication equilibrium (given the decision rule) iff

(i) $\forall i : \gamma_i^*(\sigma_i) = \begin{cases} \sigma_i, & \text{for } \sigma_i = \alpha \\ s, & \text{for } \sigma_i = \beta \end{cases}$:

$$p_i(\omega = A|\alpha_i^{piv}) \geq \theta_i, \text{ and } p_i(\omega = A|\beta_i^{piv}) > \theta_i,$$

(ii) $\forall i : \gamma_i^*(\sigma_i) = \sigma_i$:

$$p_i(\omega = A|\sigma_i = \alpha_i^{piv}) \geq \theta_i \geq p_i(\omega = A|\beta_i^{piv}), \text{ and}$$

(iii) $\forall i : \gamma_i^*(\sigma_i) = \begin{cases} s, & \text{for } \sigma_i = \alpha \\ \sigma_i, & \text{for } \sigma_i = \beta \end{cases}$:

$$p_i(\omega = A|\alpha_i^{piv}) < \theta_i, \text{ and } p_i(\omega = A|\beta_i^{piv}) \leq \theta_i,$$

where $p_i(\omega = A|\hat{\sigma}_i^{piv})$ denotes the probability i assigns to $\omega = A$ conditional on the event that his signal $\hat{\sigma}_i$ is pivotal for the decision, taking as given the communication strategies of agents $-i$ and the decision rule. We will have a closer look at $p_i(\omega = A|\hat{\sigma}^{piv})$. It is convenient to introduce some further notation.

Consider a (pure) communication strategy profile γ. Denote $\mathcal{N}_\beta(\gamma) = \{i : \gamma_i(\beta) = \beta\}$, $n_\beta(\gamma) = |\mathcal{N}_\beta(\gamma)|$, $\mathcal{N}_\alpha(\gamma) = \{i : \gamma_i(\alpha) = \alpha\}$, $n_\alpha(\gamma) = |\mathcal{N}_\alpha(\gamma)|$. Denote with $k_{-i}(c)$ the evidence provided by agents $-i$ in the communication stage.

Given communication strategies γ_{-i}, $k_{-i}(c) = k$ happens if (and only if) there are $k + l$ α-signals within the group of committee members other than i planning to reveal α, i.e. agents $j \in \mathcal{N}_\alpha(\gamma^*) \setminus \{i\}$, and l β-signals within the group of committee members (other than i) planning to reveal β (agents $j \in \mathcal{N}_\beta(\gamma^*) \setminus \{i\}$), for all $l = \max\{0, -k\}, \ldots, \min\{n_\alpha(\gamma^*) - k - 1_{\gamma_i^*(\alpha)=\alpha}, n_\beta(\gamma^*) - 1_{\gamma_i^*(\beta)=\beta}\}$, where $1_x = 1$ if x is true and 0 else. Abbreviate $L(\gamma) = \{\max\{0, -k\}, \ldots, \min\{n_\alpha(\gamma^*) - k - 1_{\gamma_i^*(\alpha)=\alpha}, n_\beta(\gamma^*) - 1_{\gamma_i^*(\beta)=\beta}\}$
An α-signal is pivotal in state A with probability

$$
\begin{aligned}
&prob\{\alpha_i^{piv}|\omega = A\} = \qquad (A.1)\\
&\sum_{l \in L(\gamma^*)} \binom{n_\alpha(\gamma^*)-1_{\gamma_i^*(\alpha)=\alpha}}{k+l} (\delta q)^{k+l} (1-\delta q)^{n_\alpha(\gamma^*)-1_{\gamma_i^*(\alpha)=\alpha}-k-l}\\
&\cdot \binom{n_\beta(\gamma^*)-1_{\gamma_i^*(\beta)=\beta}}{l} (\delta(1-q))^l (1-\delta(1-q))^{n_\beta(\gamma^*)-1_{\gamma_i^*(\beta)=\beta}-l}.
\end{aligned}
$$

In state B, an α-signal is pivotal with probability

$$
\begin{aligned}
&prob\{\alpha_i^{piv}|\omega = B\} = \qquad (A.2)\\
&\sum_{l \in L(\gamma^*)} \binom{n_\alpha(\gamma^*)-1_{\gamma_i^*(\alpha)=\alpha}}{k+l} (\delta(1-q))^{k+l} (1-\delta(1-q))^{n_\alpha(\gamma^*)-1_{\gamma_i^*(\alpha)=\alpha}-k-l}\\
&\cdot \binom{n_\beta(\gamma^*)-1_{\gamma_i^*(\beta)=\beta}}{l} (\delta q)^l (1-\delta q)^{n_\beta(\gamma^*)-1_{\gamma_i^*(\beta)=\beta}-l}.
\end{aligned}
$$

Using Bayes' Rule, we have

$$
\begin{aligned}
&p_i(\omega = A|\alpha_i^{piv}) \qquad (A.3)\\
&= \frac{p(\omega=A|\sigma_i=\alpha) prob\{\alpha_i^{piv}|\omega=A\}}{p(\omega=A|\sigma_i=\alpha) prob\{\alpha_i^{piv}|\omega=A\} + (1-p(\omega=A|\sigma_i=\alpha)) prob\{\alpha_i^{piv}|\omega=B\}}\\
&= \frac{q prob\{\alpha_i^{piv}|\omega=A\}}{q prob\{\alpha_i^{piv}|\omega=A\} + (1-q) prob\{\alpha_i^{piv}|\omega=B\}}\\
&= \frac{1}{1+\left(\frac{1-q}{q}\right)^{k+1}\left(\frac{1-\delta(1-q)}{1-\delta q}\right)^{n_\alpha(\gamma^*)-n_\beta(\gamma^*)-k-1_{\gamma_i^*(\alpha)=\alpha}+1_{\gamma_i^*(\beta)=\beta}}}.
\end{aligned}
$$

Analogously, $k_{-i}(c) = k + 1$ happens if there are $k + 1 + l$ α-signals within the group of committee members other than i planning to reveal α, i.e. agents $j \in \mathcal{N}_\alpha(\gamma^*) \setminus \{i\}$, and l β-signals within the group of committee members (other than i) planning to reveal β (agents $j \in \mathcal{N}_\beta(\gamma^*) \setminus \{i\}$), for all $l \in L(\gamma^*) \setminus \{-k\ ; n_\alpha(\gamma^*) - k - 1_{\gamma_i^*(\alpha)=\alpha}\}$. Denote $L'(\gamma^*) = L(\gamma^*) \setminus \{-k\ ; n_\alpha(\gamma^*) - k - 1_{\gamma_i^*(\alpha)=\alpha}\}$.

A β-signal is pivotal in state A with probability

$$prob\{\beta_i^{piv}|\omega = A\} = \tag{A.4}$$
$$\sum_{l \in L'(\gamma^*)} \binom{n_\alpha(\gamma^*)-1_{\gamma_i^*(\alpha)=\alpha}}{k+1+l} (\delta q)^{k+1+l} (1-\delta q)^{n_\alpha(\gamma^*)-1_{\gamma_i^*(\alpha)=\alpha}-k-1-l}$$
$$\cdot \binom{n_\beta(\gamma^*)-1_{\gamma_i^*(\beta)=\beta}}{l} (\delta(1-q))^l (1-\delta(1-q))^{n_\beta(\gamma^*)-1_{\gamma_i^*(\beta)=\beta}-l}.$$

In state B, a β-signal is pivotal with probability

$$prob\{\beta_i^{piv}|\omega = B\} = \tag{A.5}$$
$$\sum_{l \in L'(\gamma^*)} \binom{n_\beta(\gamma^*)-1_{\gamma_i^*(\beta)=\beta}}{l} (\delta q)^l (1-\delta q)^{n_\beta(\gamma^*)-1_{\gamma_i^*(\beta)=\beta}-l}$$
$$\cdot \binom{n_\alpha(\gamma^*)-1_{\gamma_i^*(\alpha)=\alpha}}{k+1+l} (\delta(1-q))^{k+1+l} (1-\delta(1-q))^{n_\alpha(\gamma^*)-1_{\gamma_i^*(\alpha)=\alpha}-k-1-l}$$

Using Bayes' Rule, we have

$$p_i(\omega = A|\beta_i^{piv}) \tag{A.6}$$
$$= \frac{p(\omega=A|\sigma_i=\beta) prob\{\beta_i^{piv}|\omega=A\}}{p(\omega=A|\sigma_i=\beta) prob\{\alpha_i^{piv}|\omega=A\} + (1-p(\omega=A|\sigma_i=\beta)) prob\{\alpha_i^{piv}|\omega=B\}}$$
$$= \frac{(1-q) prob\{\beta_i^{piv}|\omega=A\}}{(1-q) prob\{\beta_i^{piv}|\omega=A\} + (1-q) prob\{\beta_i^{piv}|\omega=B\}}$$
$$= \frac{1}{1+\left(\frac{1-q}{q}\right)^k \left(\frac{1-\delta(1-q)}{1-\delta q}\right)^{n_\alpha(\gamma^*)-n_\beta(\gamma^*)-k-1_{\gamma_i^*(\alpha)=\alpha}+1_{\gamma_i^*(\beta)=\beta}-1}}.$$

We have that $p_i(\omega = A|\beta_i^{piv}) < p_i(\omega = A|\alpha_i^{piv})$ $\forall i$. Hence, given the decision rule k, in any communication equilibrium (in pure strategies) each voter reveals at least one type of signal. Note that $p_i(\omega = A|\hat{\sigma}_i^{piv})$ is ceteris paribus higher (i) the higher k, (ii) the lower $n_\alpha(\gamma)$, and (iii) the higher $n_\beta(\gamma)$. Note also that $p_i(.)$ are the same for i (given k) for communication profiles γ' and γ'' if $n_\alpha(\gamma') - n_\beta(\gamma') = n_\alpha(\gamma'') - n_\beta(\gamma'')$ and $\gamma_i' = \gamma_i''$. Further note that $p_i(\omega = A|\hat{\sigma}_i^{piv}) = p_j(\omega = A|\hat{\sigma}_j^{piv})$ if $\gamma_i = \gamma_j$.

Consider a communication profile γ. Denote the belief $p_i(\omega = A|\alpha_i^{piv})$ of an agent who reveals both types of signals (if endowed with them), i.e. agent $i : \gamma_i(\sigma_i) = \sigma_i$ with $p_\alpha(\gamma)$ and $p_i(\omega = A|\beta_i^{piv})$ with $p_\beta(\gamma)$. Similarly, denote the beliefs of an agent who reveals α and conceals β, i.e. agent $i : \gamma_i(\alpha) = \alpha, \gamma_i(\beta) = s$ with $p_\alpha^{\beta-conc}(\gamma)$ and $p_\beta^{\beta-conc}(\gamma)$, respectively. Denote the beliefs of agent $i : \gamma_i(\alpha) = s, \gamma_i(\beta) = \beta$ with $p_\alpha^{\alpha-conc}(\gamma)$ and $p_\beta^{\alpha-conc}(\gamma)$, respectively. It is easy to verify that $p_{\hat{\sigma}}^{\beta-conc}(\gamma) > p_{\hat{\sigma}}(\gamma) > p_{\hat{\sigma}}^{\alpha-conc}(\gamma), \hat{\sigma} = \alpha, \beta$. Concerning the position of $p_\alpha^{\alpha-conc}(\gamma)$ and $p_\beta^{\beta-conc}(\gamma)$, we have to distinguish three cases:

(i) If $\delta < \delta'(q)$, we have

$$p_\beta^{\alpha-conc}(\gamma) < p_\beta(\gamma) < p_\beta^{\beta-conc}(\gamma) < p_\alpha^{\alpha-conc}(\gamma) < p_\alpha(\gamma) < p_\alpha^{\beta-conc}(\gamma),$$

(ii) if $\delta'(q) < \delta < \delta''(q)$, we have

$$p_\beta^{\alpha-conc}(\gamma) < p_\beta(\gamma) < p_\alpha^{\alpha-conc}(\gamma) < p_\beta^{\beta-conc}(\gamma) < p_\alpha(\gamma) < p_\alpha^{\beta-conc}(\gamma),$$

(iii) if $\delta > \delta''(q)$, we have

$$p_\beta^{\alpha-conc}(\gamma) < p_\alpha^{\alpha-conc}(\gamma) < p_\beta(\gamma) < p_\alpha(\gamma) < p_\beta^{\beta-conc}(\gamma) < p_\alpha^{\beta-conc}(\gamma),$$

where $\delta'(q) = \frac{1-\left(\frac{q}{1-q}\right)^{1/3}}{1-q(1+\left(\frac{q}{1-q}\right)^{1/3})}$ and $\delta''(q) = \frac{1-\left(\frac{q}{1-q}\right)^{1/2}}{1-q(1+\left(\frac{q}{1-q}\right)^{1/2})}$. The three cases are depicted in Figure A.1. In case (i), the information contained in any committee member's silence plays a minor role, because the endowment with information is relatively unlikely. Hence, beliefs are mainly determined by revealed information and the own signal. As the endowment with information becomes more likely, communication strategies of the other committee members gain importance whereas the own information endowment becomes relatively unimportant for the beliefs. As $\delta \to 1$, $p_\alpha - p_\beta \to 0$. To see why this is the case, suppose that $\delta = 1$ and suppose that n_α agents other than i reveal α and n_β agents other than i reveal β. Agent i's α is pivotal if $k + l$ α-signals and l β-signals are revealed – which implies that $(n - 1 - n_\alpha - l)$ α-signals and $(n - 1 - n_\beta - k - l)$ β-signals are concealed. Hence, i can infer that $\kappa = 2k + n_\beta - n_\alpha - 1$. He makes the same inference if he is pivotal with a β-signal. The two situations differ in that there must be an agent who has an α in the former, and a β in the latter case. As i has a β in the former and an α in the latter case, κ is inferred to be the same.

A communication profile γ is a communication equilibrium if

(i) every agent reveals at least one type of signal,

(ii) $\forall i : \theta_i < p_\alpha^{\alpha-conc}(\gamma^*) : \gamma_i^*(\alpha) = \alpha$,

(iii) $\forall i : \theta_i > p_\beta^{\beta-conc}(\gamma^*) : \gamma_i^*(\beta) = \beta$,

(iv) $\forall i : \theta_i > p_\alpha(\gamma^*) : \gamma_i^*(\alpha) = s$, and

(v) $\forall i : \theta_i < p_\beta(\gamma^*) : \gamma_i^*(\beta) = s$.

We proof existence by constructing communication profiles together with a decision rule k such that conditions (i)-(v) are met for cases (ii) and (iii). For case (i), existence is not guaranteed.

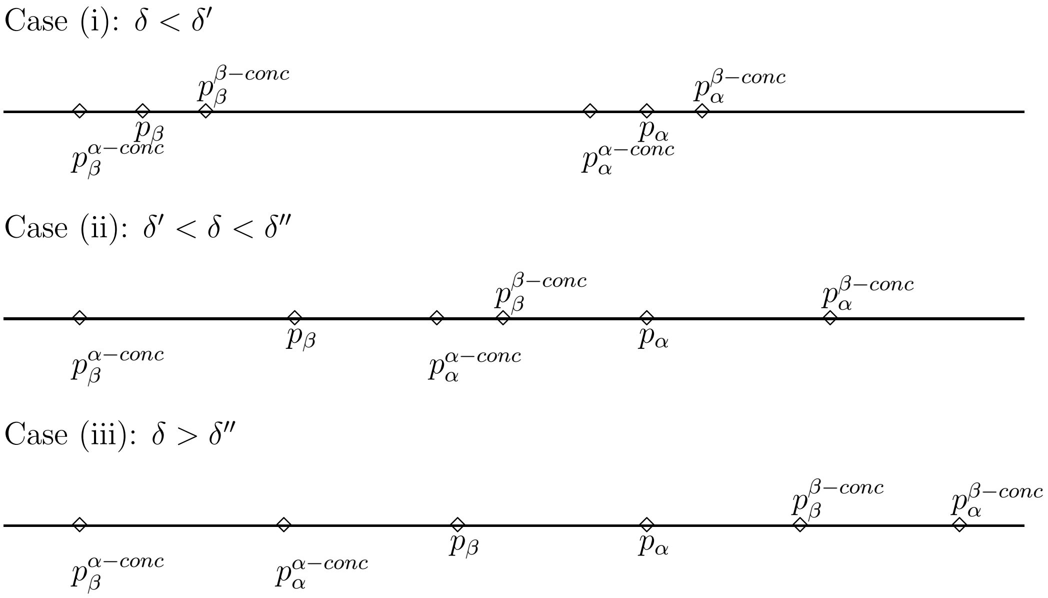

Figure A.1: Structure of the committee members' beliefs.

Consider first case (iii). Let k be the integer for which:

$$\frac{1}{1+\frac{1-\delta q}{1-\delta(1-q)}\left(\frac{(1-q)(1-\delta q)}{q(1-\delta(1-q))}\right)^k} \leq \theta_m \leq \frac{1}{1+\frac{1-q}{q}\left(\frac{(1-q)(1-\delta q)}{q(1-\delta(1-q))}\right)^k}.$$

The decision rule k is chosen in such a way that the median preference type is willing to reveal both types of signals if there are as many other agents revealing α as there are revealing β.[1] The above conditions (i)-(v) hold for the following communication profile: $\forall i : \theta_i < \theta_m : \gamma_i(\alpha) = \alpha, \gamma_i(\beta) = s$; $\forall i : \theta_i > \theta_m : \gamma_i(\alpha) = s, \gamma_i(\beta) = \beta$; $i : \theta_i = \theta_m : \gamma_i(\alpha) = \alpha, \gamma_i(\beta) = \beta$. We have $n_\alpha(\gamma) = n_\beta(\gamma) = \frac{n+1}{2}$. Hence γ is an equilibrium communication profile with the property stated in the proposition.

Consider case (ii). Let the decision rule k be such that $p_\alpha^{\alpha-conc} \leq \theta_m \leq p_\beta^{\beta-conc}$ for $n_\alpha = n_\beta$.[2] The following communication profile is an equilibrium: $\forall i : \theta_i < \theta_m : \gamma_i(\alpha) = \alpha, \gamma_i(\beta) = s$; $\forall i : \theta_i > \theta_m : \gamma_i(\alpha) = s, \gamma_i(\beta) = \beta$; $i : \theta_i = \theta_m :$

[1]As $\frac{1}{1+\frac{1-\delta q}{1-\delta(1-q)}\left(\frac{(1-q)(1-\delta q)}{q(1-\delta(1-q))}\right)^{k+1}} > \frac{1}{1+\frac{1-q}{q}\left(\frac{(1-q)(1-\delta q)}{q(1-\delta(1-q))}\right)^k}$, such a k might not exist. In this case, we can find a k together with $|n_\alpha - n_\beta| = 1$ such that the above inequalities hold. The following arguments are analogous, hence we restrict ourselves to the case that θ_m is such that a k exists for which both inequalities hold.

[2]If such a k does not exist, it exists for $|n_\alpha - n_\beta| = 1$ or 2.

$\gamma_i(\alpha) = \alpha, \gamma_i(\beta) = \beta$. Again, we have $n_\alpha(\gamma) = n_\beta(\gamma) = \frac{n+1}{2}$.

Consider case (i). Choose k such that $p_\beta^{\beta-conc} \leq \theta_m \leq p_\alpha^{\alpha-conc}$ if $n_\alpha = n_\beta$. Construct the communication profile as follows: First, let agents $i : \theta_i \leq p_\alpha^{\alpha-conc}$ reveal α and agents $i : \theta_i \geq p_\beta^{\beta-conc}$ reveal β, and let all all other information be concealed. Note that $n_\alpha, n_\beta \geq \frac{n+1}{2}$. If $n_\alpha = n_\beta$, the communication profile is an equilibrium. If $n_\alpha > n_\beta$, modify the communication profile for agents $i : p_\beta \leq \theta_i \leq p_\beta^{\beta-conc}$: let $\min\{n_\alpha - n_\beta, |\{i : p_\beta \leq \theta_i \leq p_\beta^{\beta-conc}\}|\}$ of them reveal β in addition to the revelation described above. The new communication profile is an equilibrium with the properties stated in the proposition if $n_\alpha - n_\beta \leq |\{i : p_\beta \leq \theta_i \leq p_\beta^{\beta-conc}\}|$. However, we cannot guarantee existence of such a communication equilibrium in general.

It remains to be shown that $|k| < |k_m|$. To see this, note that in equilibrium, $\theta_m \in [p_\beta, p_\alpha]$. Remember that $\theta_m \in \left[\frac{1}{1+\left(\frac{1-q}{q}\right)^{k_m}}, \frac{1}{1+\left(\frac{1-q}{q}\right)^{k_m+1}}\right]$. Given decision rule k, for $\delta \to 0$, p_β and p_α converge to $\frac{1}{1+\left(\frac{1-q}{q}\right)^{k}}$ and $\frac{1}{1+\left(\frac{1-q}{q}\right)^{k_m+1}}$, respectively. For $\delta \to 1$, they converge to $\frac{1}{1+\left(\frac{1-q}{q}\right)^{2k-n_\alpha-n_\beta+1}}$. Hence, in the equilibria which we consider $|k|$ is at most $|k_m|$ and the lower the higher δ. Q.E.D.

Probability agent i assigns to $\theta_m = \theta'$

For $\theta' < \theta_i$:

$$\phi_i^m(\theta') = (n-1)\binom{n-2}{\frac{n-1}{2}}\Phi(\theta')^{\frac{n-1}{2}}(1-\Phi(\theta'))^{\frac{n-3}{2}}\phi(\theta')d\theta', \tag{A.7}$$

for $\theta' = \theta_i$:

$$\phi_i^m(\theta_i) = \binom{n-1}{\frac{n-1}{2}}\Phi(\theta_i)^{\frac{n-1}{2}}(1-\Phi(\theta_i))^{\frac{n-1}{2}}, \tag{A.8}$$

and for $\theta' > \theta_i$:

$$\phi_i^m(\theta') = (n-1)\binom{n-2}{\frac{n-1}{2}}\Phi(\theta')^{\frac{n-3}{2}}(1-\Phi(\theta'))^{\frac{n-1}{2}}\phi(\theta')d\theta'. \tag{A.9}$$

Probability agent i assigns to κ_{-i} given σ_i

For $\sigma_i = \alpha$:

$$\mu_i(\kappa_{-i} = k'|\alpha\} = \frac{\binom{n-1}{\frac{n-1+k'}{2}}\left(\left(\frac{q}{1-q}\right)^{\frac{k'}{2}}q + \left(\frac{1-q}{q}\right)^{\frac{k'}{2}}(1-q)\right)}{\sum_{\hat{k}\in E\{-n+1,\dots n-1\}}\binom{n-1}{\frac{n-1+\hat{k}}{2}}\left(\left(\frac{q}{1-q}\right)^{\frac{\hat{k}}{2}}q + \left(\frac{1-q}{q}\right)^{\frac{\hat{k}}{2}}(1-q)\right)}, \tag{A.10}$$

and for $\sigma_i = \beta$:

$$\mu_i(\kappa_{-i} = k'|\beta) = \frac{\binom{n-1}{\frac{n-1+k'}{2}}\left(\left(\frac{q}{1-q}\right)^{\frac{k'}{2}}(1-q)+\left(\frac{1-q}{q}\right)^{\frac{k'}{2}}q\right)}{\sum\limits_{\hat{k}\in E\{-n+1,\ldots n-1\}}\binom{n-1}{\frac{n-1+\hat{k}}{2}}\left(\left(\frac{q}{1-q}\right)^{\frac{\hat{k}}{2}}q+\left(\frac{1-q}{q}\right)^{\frac{\hat{k}}{2}}(1-q)\right)}, \quad \text{(A.11)}$$

where $E\{-n+1, \ldots n-1\}$ is the set of even numbers between (including) $-n+1$ and $n-1$, and $\mu_i(\kappa_{-i} = k'|\sigma_i) = 0$ for odd values k'.

Sufficient condition for full information revelation: Example

Consider a committee with three members, each of whom receives a signal which is correct with probability 0.8. We know from Proposition 3.5 that a full information revelation equilibrium exists if the preference parameters are drawn from $[0.2, 0.985]$ or from $[0.015, 0.8]$. Suppose preferences are drawn from a uniform distribution which is symmetric with respect to 1/2. We identify the minimum θ_{min} (and therewith the maximum θ_{max}) for which a full information revelation equilibrium exists. Existence is guaranteed for $\theta_{min} \geq 0.2$, and obviously, we must have $\theta_{min} > 0.015$, otherwise the most biased types' preferred alternative does not depend on the realization of the signals. Hence, we consider the case $0.015 < \theta_{min} < 0.2$ As outlined above, out-of-equilibrium-beliefs best support a full information equilibrium, if (i) probability 1 is assigned to $\kappa = -1$ in case $c_i = s$ and $k(c) = -2$, and (ii) probability 1 is assigned to $\kappa = 3$ in case $c_i = s$ and $k(c) = 2$. Because of symmetry, out-of-equilibrium-beliefs for the case $c_i = s$ and $k(c) = 0$ assign equal probability to $\kappa = -1$ and $\kappa = 1$. Again because of symmetry, existence is guaranteed if type θ_{min} has an incentive to reveal a β-signal. Type θ_{min} has an incentive to reveal a β-signal if

$$\begin{aligned}
& (1/2 \cdot 0.8^3 + 1/2 \cdot 0.2^3)(\theta_{min} - 0.015)\Phi(0.2)(1 - \Phi(0.2)) \\
\geq \; & (1/2 \cdot 0.8 \cdot 0.2^2 + 1/2 \cdot 0.2 \cdot 0.8^2)(0.2 - \theta_{min})(\Phi(0.5) - \Phi(0.2)(1 - \Phi(0.5)) \\
\Leftrightarrow \; & 0.26(\theta_{min} - 0.015)(0.2 - \theta_{min})(0.8 - \theta_{min}) \\
> \; & 0.08(0.2 - \theta_{min})(0.5 - (0.2 - \theta_{min}))0.5 \\
\Leftrightarrow \; & \theta_{min} \geq 0.10446.
\end{aligned}$$

Note that the sufficient condition stated in Proposition 3.5 allows for potential conflicts of interests (as measured by $\theta_{max} - \theta_{min}$) of 0.785 for this example, whereas the sufficient condition stated here allows for 0.79.

Appendix B

Appendix to Chapter 4

Proof of Lemma 4.4

We have to show that neither concealing b^l nor concealing c^h is a profitable deviation for a low type from a putative equilibrium in which he reveals the information. Obviously, a low type aims to influence only the high types' votes by concealing information, as he shares common interests with the other low type. When observing c^h, a low type does not want implementation. Concealing c^h can have an effect on a high type j's vote only if j finds himself in information set $I_h(b^l)$ or $I_h(\emptyset)$. Note that $\mu_i(b^h)$ cannot be affected by concealing c^h. In both information sets, $\mu_j(c^l)$ is higher when c^h is concealed than if it is revealed. Hence, the only effect the concealment of c^h can have is to cause a vote for implementation, which is clearly not in the interest of a low type, who prefers rejection. It is easy to verify that analogous arguments apply for a low type's revelation of b^l, and a high type's revelation of b^h and c^l. Q.E.D.

Proof of Lemma 4.5

Consider an uninformed low type player i. Given that no player has revealed information, i can infer (by Lemma 4.4) that neither of the high types has observed b^h or c^l. The only player who might have information which makes the state in which i prefers implementation more likely (in the sense of raising $\mu_i(c^l)$ or $\mu_i(b^h)$ above the prior) is the other low type. If any of the high types possesses information, i can infer that either $c = c^h$ or $b = b^l$, in which case i does not want implementation. Hence, expected utility from implementation can be positive only in the case in which neither of the high types has information and the other low type has information. Agent i's vote for implementation is pivotal only if there is at least one vote against implementation, that is either if (i) the other low type votes against implementation and one or both high types vote for implementation, or if (ii) the other low type votes for implementation and at least one of the high types votes against implementation. We will show that both cases are incompatible with equilibrium behavior and the low type being informed and high types being uninformed.

Consider case (i). If expected utility from implementation for agent i is positive presuming that the other low type has information, then the other low type's expected utility is positive if he actually has the information. Hence, voting against implementation is no equilibrium action for the informed low type.

Consider case (ii). If expected utility from implementation for agent i is positive, then it must also be positive for an uninformed high type. Voting against implementation is a contradiction to equilibrium. Thus, i's vote for implementation is pivotal only in cases in which expected utility from implementation is negative.

The proof for the high types is along the same lines. Q.E.D.

Proof of Proposition 4.2

(i) In the most revealing equilibrium, the project will be implemented in state (b^h, c^l) with probability 1 if at least one committee member receives a signal b^h and at least one committee member receives a signal c^l, and with probability $\frac{1}{2}$ else. The probability that at least one committee member receives a signal b^h and at least one committee member receives a signal c^l is

$$\delta\left(1-\left(1-\tfrac{\delta}{2}\right)^3\right)+(1-\delta)\left(\delta\left(1-\left(1-\tfrac{\delta}{2}\right)^2\right)+(1-\delta)\delta\tfrac{\delta}{2}\right)$$
$$=\delta^2\left(3\left(1-\tfrac{\delta}{2}\right)^2+\tfrac{1}{2}\left(\tfrac{\delta}{2}\right)^2\right).$$

Hence, the probability of implementation is

$$\tfrac{1}{2}\left(1+\delta^2\left(3\left(1-\tfrac{\delta}{2}\right)^2+\tfrac{1}{2}\left(\tfrac{\delta}{2}\right)^2\right)\right).$$

(ii) Analogous. (iii) Obvious. Q.E.D.

Proof of Lemma 4.9

We already noticed that in a putative equilibrium with full information revelation, concealing c^l is a profitable deviation for a low type. Note that no other non-disclosure is a profitable deviation from full information revelation. It remains to be shown that concealing c^l is a profitable deviation for low type i from a putative equilibrium in which only low type j conceals c^l, and that the strategy profile consists of mutually best responses.

Consider a putative equilibrium in which low type j conceals c^l, and low type i reveals c^l, and consider i's deviation to conceal c^l. For $\frac{h-1}{h+1} < \alpha$, i does not want the project to be implemented unless he observes b^h and c^l. If b^h is revealed, high types vote for implementation, hence i can make sure implementation by voting for it. Given the putative equilibrium communication strategies, high types believe that i is uninformed if he does not reveal information. Hence, if j reveals b^l, and no information about the costs is revealed, high types assign probability $\frac{1}{2}$ to c^l and vote against implementation. They vote in favor of implementation if j reveals c^l. The deviation has a positive effect for some communication outcomes, and no negative effects. Hence, it is a profitable deviation from the putative equilibrium.

Given voting strategies, revelation strategies are mutually best responses. Low types' voting strategies obviously maximize conditional expected payoff. Note that the fact that a low type might be concealing c^l is decision-irrelevant for the other low type (Lemma 4.5).

Consider the high types' voting strategies when reaching an information set in which only b^l was revealed. If both low types reveal b^l, high types know that none of them has observed c^l. Implementation yields expected utility $\alpha - \frac{1}{2}(h-1)$, which is negative since $\alpha < \frac{h-1}{2}$. If one of the low types has not revealed any information, high types assign probability $\frac{1-\frac{\delta}{2}}{\left(1-\frac{\delta}{2}\right)+(1-\delta)}$ to $c = c^l$. Expected utility from implementation is $\frac{1-\frac{\delta}{2}}{\left(1-\frac{\delta}{2}\right)+(1-\delta)}\alpha + \frac{1-\delta}{\left(1-\frac{\delta}{2}\right)+(1-\delta)}(1+\alpha-h)$, which is negative if and only if $\delta < \frac{h-1-2\alpha}{h-1-\frac{3}{2}\alpha}$. If none of the low types reveals any information, high types assign probability $\frac{\left(1-\frac{\delta}{2}\right)^2}{\left(1-\frac{\delta}{2}\right)^2+(1-\delta)^2}$ to $c = c^l$. Expected utility from implementation is $\frac{\left(1-\frac{\delta}{2}\right)^2}{\left(1-\frac{\delta}{2}\right)^2+(1-\delta)^2}\alpha + \frac{(1-\delta)^2}{\left(1-\frac{\delta}{2}\right)^2+(1-\delta)^2}(1+\alpha-h)$, which is negative if and only if $\delta < \frac{\sqrt{h-1-\alpha}-\sqrt{\alpha}}{\sqrt{h-1-\alpha}-\frac{\sqrt{\alpha}}{2}}$. Q.E.D.

Proof of Proposition 4.3

(i) A low type votes in favor of implementation upon observing b^h and c^l in the communication stage or when observing b^h in the communication stage and (privately) observing c^l. As high types vote for implementation for all possible communication outcomes in state (b^h, c^l), one low type's vote in favor of implementation suffices in order to implement the project with probability 1. Hence, the same distribution of information yields implementation with certainty as for $\alpha > \frac{h-1}{2}$. The proof of Proposition 4.2(i) applies.

(ii) Concerning the probability of rejection in state (b^l, c^h) we have to distinguish three cases. Low types vote against implementation for any communication outcome. If $\delta > \frac{h-1-2\alpha}{h-1-\frac{3}{2}\alpha}$, high types vote against implementation only if the state is revealed or both low types revealed b^l. The probability of these events is

$$\delta\left(\tfrac{1}{2}\left(\delta + (1-\delta)\left(1-\left(1-\tfrac{\delta}{2}\right)^2\right)\right) + \tfrac{1}{2}\left(1-\left(1-\tfrac{\delta}{2}\right)^3\right)\right)$$
$$+(1-\delta)\left(\delta\left(1-\left(1-\tfrac{\delta}{2}\right)^2\right) + (1-\delta)2\left(\tfrac{\delta}{2}\right)^2\right)$$
$$= \delta^2\left(\tfrac{13}{4}\left(1-\tfrac{\delta}{2}\right)^2 + \tfrac{1}{2}\left(\tfrac{\delta}{2}\right)^2\right).$$

Hence, the probability of rejection is

$$\tfrac{1}{2}\left(1+\delta^2\left(\tfrac{13}{4}\left(1-\tfrac{\delta}{2}\right)^2 + \tfrac{1}{2}\left(\tfrac{\delta}{2}\right)^2\right)\right).$$

If $\frac{\sqrt{h-1-\alpha}-\sqrt{\alpha}}{\sqrt{h-1-\alpha}-\frac{\sqrt{\alpha}}{2}} < \delta \leq \frac{h-1-2\alpha}{h-1-\frac{3}{2}\alpha}$, high types vote against implementation if they learn the state of the world or if they learn b^l and at least one of the low types reveals information. The probability of these events is

$$\delta\left(\tfrac{1}{2} + \tfrac{1}{2}\left(1-\left(1-\tfrac{\delta}{2}\right)^3\right)\right) + (1-\delta)\left(\delta\left(\tfrac{1}{2}+\tfrac{1}{2}\left(1-\left(1-\tfrac{\delta}{2}\right)^2\right)\right) + (1-\delta)2\left(\tfrac{\delta}{2}\right)^2\right)$$
$$= \delta\left(1+\tfrac{5}{4}\delta - 2\delta^2 + \tfrac{11}{16}\delta^3\right).$$

Hence, the probability of rejection is

$$\tfrac{1}{2}\left(1+\delta\left(1+\tfrac{5}{4}\delta - 2\delta^2 + \tfrac{11}{16}\delta^3\right)\right).$$

If $\delta \leq \frac{\sqrt{h-1-\alpha}-\sqrt{\alpha}}{\sqrt{h-1-\alpha}-\frac{\sqrt{\alpha}}{2}}$, high types vote against implementation if they learn b^l. The probability of this event is

$$1-\left(1-\left(\tfrac{\delta}{2}\right)\right)^4.$$

Hence, the probability of rejection is

$$1-\tfrac{1}{2}\left(1-\left(\tfrac{\delta}{2}\right)\right)^4.$$

(iii) Obvious. Q.E.D.

Proof of Proposition 4.4

(i) We have to show that

(a) $1-\frac{1}{2}\left(1-\frac{\delta}{2}\right)^2-\frac{1}{2}\left(1+\delta^2\left(\frac{13}{4}\left(1-\frac{\delta}{2}\right)^2+\frac{1}{2}\left(\frac{\delta}{2}\right)^2\right)\right)>0$ for $\delta<\delta^*$, and

(b) $\delta^*>\frac{h-1-2\alpha}{h-1-\frac{3}{2}\alpha}$.

(a) $1-\frac{1}{2}\left(1-\frac{\delta}{2}\right)^2-\frac{1}{2}\left(1+\delta^2\left(\frac{13}{4}\left(1-\frac{\delta}{2}\right)^2+\frac{1}{2}\left(\frac{\delta}{2}\right)^2\right)\right)$

$=\frac{\delta}{2}\left(1-\frac{7}{2}\delta+\frac{13}{4}\delta^2-\frac{15}{16}\delta^3\right)$, which is positive if $1-\frac{7}{2}\delta+\frac{13}{4}\delta^2-\frac{15}{16}\delta^3>0$. It is easy to verify that the second factor monotonously decreases in δ and is positive for $\delta=\frac{4}{9}$.

(b) $\frac{h-1-2\alpha}{h-1-\frac{3}{2}\alpha}$ monotonously decreases in α. For the smallest value in the parameter range, $\alpha=\frac{h-1}{h+1}$, we have $\frac{h-1-2\alpha}{h-1-\frac{3}{2}\alpha}=\frac{2}{3}(h-1)$, which is smaller than $\frac{4}{9}$ for $h<\frac{5}{3}$. Hence, existence of a parameter range for δ for which the party system performs better than the parliament is guaranteed for the entire relevant parameter range if $h<\frac{5}{3}$.

(ii) We have to show that

$(h-1)\frac{1}{2}\left(\delta+\left(\frac{\delta}{2}\right)^2\right)>(h-1)\frac{1}{2}\delta^2\left(\left(\frac{\delta}{2}\right)^2+\frac{25}{4}\left(1-\frac{\delta}{2}\right)^2\right)$.

$$\begin{aligned}
\frac{1}{4}(h-1)\frac{1}{2}\left(\delta+\left(\frac{\delta}{2}\right)^2\right) &> \frac{1}{4}(h-1)\frac{1}{2}\delta^2\left(\left(\frac{\delta}{2}\right)^2+\frac{25}{4}\left(1-\frac{\delta}{2}\right)^2\right)\\
\Leftrightarrow \delta+\left(\frac{\delta}{2}\right)^2 &> \delta^2\left(\left(\frac{\delta}{2}\right)^2+\frac{25}{4}\left(1-\frac{\delta}{2}\right)^2\right)\\
\Leftrightarrow 1 &> \delta\left(\left(\frac{\delta}{2}\right)^2+\frac{25}{4}\left(1-\frac{\delta}{2}\right)^2-\frac{1}{4}\right).
\end{aligned}$$

It is easy to verify that the right-hand-side is smaller than 1 for $\delta=\frac{1}{5}$, and monotonously increasing for smaller values. Existence of the parameter ranges is obvious, as $\frac{h-1-2\alpha}{h-1-\frac{3}{2}\alpha}$ is zero for $\alpha=\frac{h-1}{2}$. Q.E.D.

Bibliography

[1] Adler, Stephan J. (1994): *The Jury: Trial and Error in the American Courtroom*, New York: Times Books.

[2] Aghion, Philippe and Jean Tirole (1997): Formal and Real Authority in Organizations, *Journal of Political Economy* 105(1):1-29.

[3] Austen-Smith, David and Jeffrey S. Banks (1996): "Information Aggregation, Rationality, and the Condorcet Jury Theorem", *American Political Science Review* 90(1):34-45.

[4] Austen-Smith, David and Timothy Feddersen (2002): "Deliberation and Voting Rules", mimeo.

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

[6] Austen-Smith, David (1990b): "Credible Debate Equilibria", *Social Choice and Welfare* 7:75-93.

[7] Austen-Smith, David and Timothy Feddersen (2006): "Deliberation, Preference Uncertainty, and Voting Rules", *American Political Science Review* 100(2):209-217.

[8] Banerjee, Abhijit and Rohini Somanathan (2001): "A Simple Model of Voice", *Quarterly Journal of Economics* 116:189-228.

[9] Beniers, Klaas and Otto Swank (2004) "On the Composition of Committees", *Journal of Law, Economics and Organization* 20:353-378.

[10] Ben-Yashar, Ruth, Shmuel Nitzan (1997): "The optimal decision rule for fixed size committees in dichotomous choice situations: The general result." *International Economic Review* 38(1):175-187

[11] Ben-Yashar, Ruth and Smuhel Nitzan (2001): "The Invalidity of the Condorcet Jury Theorem under Endogenous Decision Skills", *Economics of Governance* 2:243-249.

[12] Ben-Yashar, Ruth and Jacob Paroush (2001): "Optimal decision rules for fixed-size committees in polychotomous choice situations", *Social Choice and Welfare* 18:737-746.

[13] Blinder, Alan S. and John Morgan (2000) "Are two heads better than one?: An experimental analyses of group versus individual decision making", NBER Working Paper 7909.

[14] Bolton, Patrick and Mathias Dewatripont (1994): "The Firm as a Communicating Network", *Quarterly Journal of Economics* 109:809-839.

[15] Cai, Hongbin (2001): "Optimal Committee Design with Heterogeneous Preferences", UCLA, Working Paper.

[16] Chwe, Michael Suk-Young (1999): "Minority Voting Rights Can Maximize Majority Welfare", *American Political Science Review* 93:65-97.

[17] Chwe, Michael Suk-Young (2006): "A Robust and Optimal Procedure for Condorcet's Model", mimeo.

[18] Condorcet, Marquis de (1785): *Essai sur l'application de l'analyse à la probabilité des decisions rendues a la pluralité des voix*, Paris: L'imprimerie royale.

[19] Costinot, Arnaud and Navin Kartik (2006): "Information Aggregation, Strategic Voting, and Institutional Design", mimeo.

[20] Coughlan, Pete (2000): "In Defence of Unanimous Jury Verdicts: Communication, Mistrials, and Sincerity", *American Political Science Review* 94:375-393.

[21] Coupé, Tom and Abdul G. Noury (2002): "On Choosing Not To Choose: Testing The Swing Voter's Curse", ECARES, Université Libre de Bruxelles, Working Paper.

[22] Crémer, Jacques (1980): A Partial Theory of the Optimal Organization of a Bureaucracy, *Bell Journal of Economics* 11:683-93.

[23] Dal Bo, Ernesto (2003): "Bribing voters", manuscript, Berkeley University.

[24] DeGroot (1970): *Optimal Statistical Decisions*, McGraw-Hill Book Company, New York.

[25] Dewatripont, Mathias and Jean Tirole (1999): "Advocates", *Journal of Political Economy* 107:1-39.

[26] Dewatripont, Mathias and Jean Tirole (2004): Modes of Communication, IDEI Working Paper 323.

[27] Doraszelski, Ulrich, Dino Gerardi and Francesco Squintani (2003): "Communication and Voting with Double-Sided Information", *Contributions to Theoretical Economics*, 3(1), Article 6.

[28] Duggan, John and César Martinelli (2001): "A Bayesian Model of Voting in Juries", *Games and Economic Behavior* 37:259-294.

[29] Dur, Robert and Otto Swank (2005): "Producing and Manipulating Information", *Economic Journal* 115:185-199.

[30] Faust, Jon (1996): "Whom can we trust to run the Fed? Theoretical support for the founder's view", *Journal of Monetary Economics* 37:267-283.

[31] Felgenhauer, Mike and Hans Peter Grüner (2003): "Committees and Special Interests", ECB working paper No. 293.

[32] Feddersen, Timothy J. and Wolfgang Pesendorfer (1996): "The Swing Voter's Curse", *American Economic Review* 86(3):408-424.

[33] Feddersen, Timothy J. and Wolfgang Pesendorfer (1997): "Voting Behavior and Information Aggregation in Elections with Private Information", *Econometrica*, 65(5):1029-1058.

[34] Feddersen, Timothy J. and Wolfgang Pesendorfer (1998): "Convicting the Innocent: The Inferiority of Unanimous Jury Verdicts under Strategic Voting", *American Political Science Review* 92(1):23-35.

[35] Feddersen, Timothy J. and Wolfgang Pesendorfer (1999a): "Election, Information Aggregation and Strategic Voting", *Proceedings of the National Academy of Sciences* 96:10572-10574.

[36] Feddersen, Timothy J. and Wolfgang Pesendorfer (1999b): "Abstention in Elections with Asymmetric Information and Diverse Preferences", *American Political Science Review* 93(2):381-398.

[37] Fey, Mark and Jaehoon Kim (2002): "The Swing Voter's Curse: Comment", *American Economic Review* 92(4):1264-1268.

[38] Garicano, Luis (2000): Hierarchies and the Organization of Knowledge in Production, *Journal of Political Economy* 108(5):874-904.

[39] Gerardi, Dino (2000): "Jury Verdicts and Preference Diversity", *American Political Science Review* 94:395-406.

[40] Gerardi, Dino, Richard McLean and Andrew Postlewaite (2005): "Aggregation of Expert Opinions", Cowles Foundation Discussion Paper 1503.

[41] Gerardi, Dino and Leeat Yariv (2003a): "Putting your Ballot where your Mouth is an Analysis of Collective Choice with Communication", UCLA Working Paper 827.

[42] Gerardi, Dino and Leeat Yariv (2003b): "Committee Design in the Presence of Communication", Cowles Foundation Discussion Paper 1411.

[43] Gerardi, Dino and Leeat Yariv (2005): "Deliberative Voting", forthcoming *Journal of Economic Theory*

[44] Gerlach-Kristen, Petra (2002): "Monetary Policy Committees and Interest Rate Setting", forthcoming in *European Economic Review.*

[45] Gerling, Kerstin, Hans Peter Grüner, Alexandra Kiel and Elisabeth Schulte (2005): "Information Acquisition and Decision Making in Committees: a Survey", *European Journal of Political Economy* 21:563-597.

[46] Gersbach, Hans and Volker Hahn (2000): "Should the Individual Voting Records of Central Bankers be Published?", mimeo, University of Heidelberg.

[47] Gershkov, Alex and Balász Szentes (2005): "Optimal Voting Schemes with Costly Information Acquisition", mimeo.

[48] Gilligan, Thomas W. and Keith Krehbiel (1987): "Collective Decision Making and Standing Committees: An Information Rationale for Restrictive Amendment Procedures", *Journal of Law, Economics and Organisation* 3:287-335.

[49] Gibbons, Allen and Wojciech Rytter (1988): *Efficient Parallel Algorithms*, Cambridge: Cambridge University Press.

[50] Grüner, Hans Peter and Elisabeth Schulte (2004): "Speed and Quality of Collective Decision Making, II: Incentives for Information Provision", CEPR Discussion Paper 4397.

[51] Grüner, Hans Peter and Alexandra Kiel (2004): "Collective Decisions with Interdependent Valuations", *European Economic Review* 48:1147-1168.

[52] Guarnaschelli, Serena, Richard D. McKelvey and Thomas R. Palfrey (2000): "Experimental Study of Jury Decision Rules", *American Political Science Review* 94:407-423.

[53] Jehiel, Philippe (1999): "Information Aggregation and Communication in Organizations", *Management Science* 45:659-669.

[54] Karotkin, Dora and Shmuel Nitzan (1995): "Two Remarks on the Effect of Increased Equalitarianism in Decisional Skills on the Number of Individuals that Maximizes Group Judgmental Competence", *Public Choice* 85:307-311.

[55] Klaven, Harry and Hans Zeisel (1966): *The American Jury*, Boston: Little, Brown.

[56] Koh, Winston TH (1994): "Making decisions in committees: A human fallibility approach." *Journal of Economic Behavior and Organization* 23:195-214.

[57] Koh, Winston TH (2004a) Optimal Sequential Decision Architectures and the Robustness of Hierarchies and Polyarchies", forthcoming *Social Choice and Welfare*.

[58] Koh, Winston TH (2004b) "The Optimal Design of Fallible Organizations: Invariance of Optimal Decision Criterion and Uniqueness of Hierarchy and Polyarchy Structures", forthcoming *Social Choice and Welfare*.

[59] Li, Hao, Sherwin Rosen and Wing Suen (2001): "Conflicts and Common Interests in Committees", *American Economic Review* 91(5):1478-1497.

[60] Lohmann, Susanne (1992): "Optimal Commitment in Monetary Policy: Credibility versus Flexibility", *American Economic Review* 82(1):273-86.

[61] Luce, R. Duncan (1959): *Individual Choice Behavior*. New York: Wiley.

[62] Martinelli, César (2002a): "Would Rational Voters Acquire Costly Information?", Instituto Technológico Autónomo de México, Discussion Paper 02-10.

[63] Martinelli, César (2002b): "Convergence Results for Unanimous Voting", *Journal of Economic Theory* 105, 278-297.

[64] Maug, Ernst and Bilge Yilmaz (2002): "Two-Class Voting: A Mechanism for Conflict Resolution", *American Economic Review* 92(5):1448-1471.

[65] McLennan, Andrew (1998): "Consequences of the Condorcet Jury Theorem for Beneficial Information Aggregation by Rational Agents" *American Political Science Review* 92:413-418.

[66] Meagher, Kieron, Hakan Orbay and Timothy Van Zandt (2001): "Hierarchy Size and Environmental Uncertainty" CEPR Discussion Paper No. 2839.

[67] Meagher, Kieron and Timothy Van Zandt (1998): Managerial Costs for One-shot Decentralized Information Processing, *Review of Economic Design* 3:329-345.

[68] Meirowitz, Adam (2005): "Designing Institutions to Aggregate Preferences and Information", Princeton University, mimeo.

[69] Melumad, Nahum D., Dilip Mookherjee and Stefan Reichelstein (1995): "Hierarchical Decentralization of Incentive Contracts", *RAND Journal of Economics* 26:654-672.

[70] Milgrom, Paul (1979): "A Convergence Theorem for Competitive Bidding with Differential Information", *Econometrica* 47:679-688.

[71] Milgrom, Paul and John Roberts (1986): "Relying on the Information of Interested Parties", *RAND Journal of Economics* 17:18-32.

[72] Mookherjee, Dilip and Stefan Reichelstein (1997): "Budgeting and Hierarchical Control", *Journal of Accounting Research* 35(2):129-155.

[73] Mukhopadhaya, Kaushik (2003): "Jury Size and the Free Rider Problem", *Journal of Law, Economics and Organization* 19:24-44.

[74] Nitzan, Smuhel and Jacob Paroush (1982): "Optimal Decision Rules in Uncertain Dichotomous Choice Situations", *International Economic Review* 23:257-297.

[75] Orbay, Hakan (2002): "Information Processing Hierarchies", *Journal of Economic Theory* 105: 370-407.

[76] Ottaviani, Marco, and Peter Sorensen (2001): "Information aggregation in debate: who should speak first", *Journal of Public Economics* 81, 393-421.

[77] Persico, Nicola (2004): "Committee Design with Endogenous Information", *Review of Economic Studies* 71:165-191.

[78] Piketty, Thomas (1999) "The Information Aggregation Approach to Political Institutions" *European Economic Review* 43:791-800.

[79] Prat, Andrea (1997): "Hierarchies of Processors with Endogenous Capacity", *Journal of Economic Theory* 77: 214-222.

[80] Radner, Roy (1992): "Hierarchy: The Economics of Managing", *Journal of Economic Literature* 30: 1382-1415.

[81] Radner, Roy (1993): "The Organization of Decentralized Information Processing", *Econometrica* 61(5): 1109-1146.

[82] Radner, Roy and Timothy Van Zandt (1992): "Information Processing in Firms and Returns to Scale", *Annales d'economie et de statistique* 25/26:265-298.

[83] Rosen, Sherwin (1982): "Autority, control, and the distribution of earnings", *Bell Journal of Economics* 13:311-323.

[84] Sah, Raaj K. and Joseph E. Stiglitz (1986): "The Architecture of Economic Systems: Hierarchies and Polyarchies", *The American Economic Review* 76(4):716-727.

[85] Sah, Raaj K. and Joseph E. Stiglitz (1988): "Committees, hierarchies and polyarchies", *Economic Journal* 98:451-470.

[86] Schulte, Elisabeth and Hans Peter Grüner (2004): "Speed and Quality of Collective Decision Making: Imperfect Information Processing", forthcoming *Journal of Economic Theory*.

[87] Schulte, Elisabeth and Hans Peter Grüner (2004): "Speed and Quality of Collective Decision Making, I: Imperfect Information Processing", CEPR Discussion Paper 4179.

[88] Seidmann, Daniel, (2004) "A Theory of Voting Patterns and Performance in Private and Public Committees", mimeo, Nottingham University.

[89] Sibert, Anne and Ilian Mihov (2002): "Credibility and Flexibility with Monetary Policy Committees", CEPR Discussion Paper 3278.

[90] Smorodinsky, Rann and Mose Tenneholtz (2003): "Overcoming Free Riding in Multi-Party Computations – The Anonymous Case", Technion, mimeo.

[91] Van Zandt, Timothy (1997): "The Scheduling and Organization of Periodic Associative Computation: Essential networks", *Review of Economic Design* 3:15-27.

[92] Van Zandt, Timothy (1998): "The Scheduling and Organization of Periodic Associative Computation: Efficient networks", *Review of Economic Design* 3:93-127.

[93] Van Zandt, Timothy (2003): "Balancedness of Real-Time Hierarchical Resource Allocation", INSEAD.

[94] Winter, Eyal (2004): "Incentives and Discrimination", *American Economic Review* 94:764-773.

[95] Wolinsky, Asher (2002): "Eliciting Information from Multiple Experts", *Games and Economic Behavior* 41(1):141-160.

Curriculum Vitae

Personal Details

Name	Elisabeth Schulte
Born	09.05.1976 in Haren (Ems), Germany

Education

2001-2006	Ph.D. studies at Mannheim University Research and teaching assistant at the chair for economic policy, Mannheim University
2004-2005	Marie Curie fellow at Université Toulouse I, France
2003-2004	Member of the interdisciplinary research group "Institutionalization of International Negotiation Systems"
1995-2001	Undergraduate studies in Business Administration, University of Osnabrück
1989-1995	Gymnasium Haren (Ems)